

RICHARD T. ELY LECTURE

The Economics of Resources or the Resources of Economics

By ROBERT M. SOLOW*

It is easy to choose a subject for a distinguished lecture like this, before a large and critical audience with a wide range of interests. You need a topic that is absolutely contemporary, but somehow perennial. It should survey a broad field, without being superficial or vague. It should probably bear some relation to economic policy, but of course it must have some serious analytical foundations. It is nice if the topic has an important literature in the past of our subject—a literature which you can summarize brilliantly in about eleven minutes—but it better be something in which economists are interested today, and it should appropriately be a subject you have worked on yourself. The lecture should have some technical interest, because you can't waffle for a whole hour to a room full of professionals, but it is hardly the occasion to use a blackboard.

I said that it is easy to choose a subject for the Ely Lecture. It has to be, because twelve people, counting me, have done it.

I am going to begin with a quotation that could have come from yesterday's newspaper, or the most recent issue of the *American Economic Review*.

Contemplation of the world's disappearing supplies of minerals, forests, and other exhaustible assets has led to demands for regulation of their exploitation. The feeling that these products are now too cheap for the good of future generations, that they are being selfishly

exploited at too rapid a rate, and that in consequence of their excessive cheapness they are being produced and consumed wastefully has given rise to the conservation movement.

The author of those sentences is not Dennis Meadows and associates, not Ralph Nader and associates, not the President of the Sierra Club; it is a very eminent economic theorist, a Distinguished Fellow of this Association, Harold Hotelling, who died at the age of seventy-eight, just a few days ago. Like all economic theorists, I am much in his debt, and I would be happy to have this lecture stand as a tribute to him. These sentences appeared at the beginning of his article "The Economics of Exhaustible Resources," not in the most recent *Review*, but in the *Journal of Political Economy* for April 1931. So I think I have found something that is both contemporary and perennial. The world has been exhausting its exhaustible resources since the first cave-man chipped a flint, and I imagine the process will go on for a long, long time.

Mr. Dooley noticed that "th' Supreme Court follows the illiction returns." He would be glad to know that economic theorists read the newspapers. About a year ago, having seen several of those respectable committee reports on the advancing scarcity of materials in the United States and the world, and having, like everyone else, been suckered into reading the *Limits to Growth*, I decided I ought to find out what economic theory has to say

* Professor of economics, Massachusetts Institute of Technology.

about the problems connected with exhaustible resources. I read some of the literature, including Hotelling's classic article—the theoretical literature on exhaustible resources is, fortunately, not very large—and began doing some work of my own on the problem of optimal social management of a stock of a nonrenewable but essential resource. I will be mentioning some of the results later. About the time I finished a first draft of my own paper and was patting myself on the back for having been clever enough to realize that there was in fact something still to be said on this important, contemporary but somehow perennial topic—just about then it seemed that every time the mail came it contained another paper by another economic theorist on the economics of exhaustible resources.¹ It was a little like trotting down to the sea, minding your own business like any nice independent rat, and then looking around and suddenly discovering that you're a lemming. Anyhow, I now have a nice collection of papers on the theory of exhaustible resources; and most of them are still unpublished, which is just the advantage I need over the rest of you.

A pool of oil or vein of iron or deposit of copper in the ground is a capital asset to society and to its owner (in the kind of society in which such things have private owners) much like a printing press or a building or any other reproducible capital asset. The only difference is that the natural resource is not reproducible, so the size of the existing stock can never in-

crease through time. It can only decrease (or, if none is mined for a while, stay the same). This is true even of recyclable materials; the laws of thermodynamics and life guarantee that we will never recover a whole pound of secondary copper from a pound of primary copper in use, or a whole pound of tertiary copper from a pound of secondary copper in use. There is leakage at every round; and a formula just like the ordinary multiplier formula tells us how much copper use can be built on the world's initial endowment of copper, in terms of the recycling or recovery ratio. There is always less ultimate copper use left than there was last year, less by the amount dissipated beyond recovery during the year. So copper remains an exhaustible resource, despite the possibility of partial recycling.

A resource deposit draws its market value, ultimately, from the prospect of extraction and sale. In the meanwhile, its owner, like the owner of every capital asset, is asking: What have you done for me lately? The only way that a resource deposit in the ground and left in the ground can produce a current return for its owner is by appreciating in value. Asset markets can be in equilibrium only when all assets in a given risk class earn the same rate of return, partly as current dividend and partly as capital gain. The common rate of return is the interest rate for that risk class. Since resource deposits have the peculiar property that they yield no dividend so long as they stay in the ground, in equilibrium the value of a resource deposit must be growing at a rate equal to the rate of interest. Since the value of a deposit is also the present value of future sales from it, after deduction of extraction costs, resource owners must expect the net price of the ore to be increasing exponentially at a rate equal to the rate of interest. If the mining industry is competitive, net price stands for market price minus margi-

¹ *The Review of Economic Studies* will publish a group of them in the summer of 1974, including my own paper and others by Partha Dasgupta and Geoffrey Heal, Michael Weinstein and Richard Zeckhauser, and Joseph Stiglitz, from all of which I have learned a lot about this subject. I would especially like to thank Zeckhauser for conversation and correspondence, and for the kind of reading of the first draft of this lecture that one only dares to hope to get because it is so close to Christmas. The final version reflects his comments.

nal extraction cost for a ton of ore. If the industry operates under constant costs, that is just market price net of unit extraction costs, or the profit margin. If the industry is more or less monopolistic, as is frequently the case in extractive industry, it is the marginal profit—marginal revenue less marginal cost—that has to be growing, and expected to grow, proportionally like the rate of interest.

This is the fundamental principle of the economics of exhaustible resources. It was the basis of Hotelling's classic article. I have deduced it as a condition of stock equilibrium in the asset market. Hotelling thought of it mainly as a condition of flow equilibrium in the market for ore: if net price is increasing like compound interest, owners of operating mines will be indifferent at the margin between extracting and holding at every instant of time. So one can imagine production just equal to demand at the current price, and the ore market clears. No other time profile for prices can elicit positive production in every period of time.

It is hard to overemphasize the importance of this tilt in the time profile for net price. If the net price were to rise too slowly, production would be pushed nearer in time and the resource would be exhausted quickly, precisely because no one would wish to hold resources in the ground and earn less than the going rate of return. If the net price were to rise too fast, resource deposits would be an excellent way to hold wealth, and owners would delay production while they enjoyed supernormal capital gains.

According to the fundamental principle, if we observe the market for an exhaustible resource near equilibrium, we should see the net price—or marginal profit—rising exponentially. That is not quite the same thing as seeing the market price to users of the resource rising exponentially. The price to consumers is the net price plus

extraction costs, or the obvious analogy for monopoly. The market price can fall or stay constant while the net price is rising if extraction costs are falling through time, and if the net price or scarcity rent is not too large a proportion of the market price. That is presumably what has been happening in the market for most exhaustible resources in the past. (It is odd that there are not some econometric studies designed to find out just this. Maybe econometricians don't follow the illicion returns.) Eventually, as the extraction cost falls and the net price rises, the scarcity rent must come to dominate the movement of market price, so the market price will eventually rise, although that may take a very long time to happen. Whatever the pattern, the market price and the rate of extraction are connected by the demand curve for the resource. So, ultimately, when the market price rises, the current rate of production must fall along the demand curve. Sooner or later, the market price will get high enough to choke off the demand entirely. At that moment production falls to zero. If flows and stocks have been beautifully coordinated through the operations of futures markets or a planning board, the last ton produced will also be the last ton in the ground. The resource will be exhausted at the instant that it has priced itself out of the market. The Age of Oil or Zinc or Whatever It Is will have come to an end. (There is a limiting case, of course, in which demand goes asymptotically to zero as the price rises to infinity, and the resource is exhausted only asymptotically. But it is neither believable nor important.)

Now let us do an exercise with this apparatus. Suppose there are two sources of the same ore, one high-cost and the other low-cost. The cost difference may reflect geographical accessibility and transportation costs, or some geological or chemical difference that makes extraction cheap at one

site and dear at the other. The important thing is that there are cost differences, though the final mineral product is identical from both sources.

It is easy to see that production from both sources cannot coexist in the market for any interval of time. For both sources to produce, net price for each of them must be growing like compound interest at the market rate. But they must market their ore at the same price, because the product is identical. That is arithmetically impossible, if their extraction costs differ.

So the story has to go like this. First one source operates and supplies the whole market. Its net price rises exponentially, and the market price moves correspondingly. At a certain moment, the first source is exhausted. At just that moment and not before, it must become economical for the second source to come into production. From then on, the world is in the single-source situation: the net price calculated with current extraction costs must rise exponentially until all production is choked off and the second source is exhausted. (If there are many sources, you can see how it will work.)

Which source will be used first? Your instinct tells you that the low-cost deposit will be the first one worked, and your instinct is right. You can see why, in terms of the fundamental principle. At the beginning, if the high-cost producer is serving the market, the market price must cover high extraction costs plus a scarcity rent that is growing exponentially. The low-cost producer would refrain from undercutting the price and entering the market only if his capital gains justify holding off and entering the market later. But just the reverse will be true. Any price high enough to keep the high-cost producer in business will tempt the low-cost producer to sell ore while the selling is good and invest the proceeds in any asset paying the market rate of interest. So it

must be that the low-cost producer is the first to enter. Price rises and output falls. Eventually, at precisely the moment when the low-cost supply is exhausted, the price has reached a level at which it pays the high-cost producer to enter. From then on, his net price rises exponentially and production continues to fall. When cumulative production has exhausted the high-cost deposit, the market price must be such as to choke the demand off to zero—or else just high enough to tempt a still higher-cost source into production. And so it goes. Apart from market processes, it is actually socially rational to use the lower-cost deposits before the higher-cost ones.

You can take this story even further, as William Nordhaus has done in connection with the energy industry. Suppose that, somewhere in the background, there is a technology capable of producing or substituting for a mineral resource at relatively high cost, but on an effectively inexhaustible resource base. Nordhaus calls this a "backstop technology." (The nearest we now have to such a thing is the breeder reactor using U^{233} as fuel. World reserves of U^{233} are thought to be enough to provide energy for over a million years at current rates of consumption. If that is not a backstop technology, it is at least a catcher who will not allow a lot of passed balls. For a better approximation, we must wait for controlled nuclear fusion or direct use of solar energy. The sun will not last forever, but it will last at least as long as we do, more or less by definition.) Since there is no scarcity rent to grow exponentially, the backstop technology can operate as soon as the market price rises enough to cover its extraction costs (including, of course, profit on the capital equipment involved in production). And as soon as that happens, the market price of the ore or its substitute stops rising. The "backstop technology" provides a ceiling for the market price of the natural resource.

The story in the early stages is as I have told it. In the beginning, the successive grades of the resource are mined. The last and highest-cost source gives out just when the market price has risen to the point where the backstop technology becomes competitive. During the earlier phases, one imagines that resource companies keep a careful eye on the prospective costs associated with the backstop technology. Any laboratory success or failure that changes those prospective costs has instantaneous effects on the capital value of existing resource deposits, and on the most profitable rate of current production. In actual fact, those future costs have to be regarded as uncertain. A correct theory of market behavior and a correct theory of optimal social policy will have to take account of technological uncertainty (and perhaps also uncertainty about the true size of mineral reserves).

Here is a mildly concrete illustration of these principles. There is now a workable technology for liquefying coal—that is, for producing synthetic crude oil from coal.¹ Nordhaus puts the extraction-and-preparation cost at the equivalent of seven or eight 1970 dollars per barrel of crude oil, including amortization and interest at 10 percent on the plant; I have heard higher and lower figures quoted. If coal were available in unlimited amounts, that would be all. But, of course, coal is a scarce resource, though more abundant than drillable petroleum, so a scarcity rent has to be added to that figure, and the rent has to be increasing like the rate of interest during the period when coal is being used for this purpose.

In the meanwhile, the extraction and production cost for this technology is

large compared with the scarcity rent on the coal input, so the market price at which the liquefied-coal-synthetic-crude activity would now be economic is rising more slowly than the rate of interest. It may even fall if there are cost-reducing technological improvements; and that is not unlikely, given that research on coal has not been splashed as liberally with funds as research on nuclear energy. In any case, political shenanigans and monopoly profits aside, scarcity rents on oil form a larger fraction of the market price of oil, precisely because it is a lower cost fuel. The price of a barrel of oil should therefore be rising faster than the implicit price at which synthetic crude from coal could compete. One day those curves will intersect, and that day the synthetic-crude technology will replace the drilled-petroleum technology.

Even before that day, the possibility of coal liquefaction provides a kind of ceiling for the price of oil. I say "kind of" to remind you that coal-mining and moving capacity and synthetic-crude plant cannot be created overnight. One might hope that the ceiling might also limit the consuming world's vulnerability to political shenanigans and monopoly profits. I suppose it does in some ultimate sense, but one must not slide over the difficulties: for example, who would want to make a large investment in coal liquefaction or coal gasification in the knowledge that the current price of oil contains a large monopoly element that could be cut, at least temporarily, if something like a price war should develop?

The fundamental principle of the economics of exhaustible resources is, as I have said, simultaneously a condition of flow equilibrium in the market for the ore and of asset equilibrium in the market for deposits. When it holds, it says quite a lot about the probable pattern of exploitation of a resource. But there are more than the usual reasons for wondering whether

¹ As best one can tell at the moment, shale oil is a more likely successor to oil and natural gas than either gasified or liquefied coal. The relevant costs are bound to be uncertain until more research and development has been done. I tell the story in terms of liquefied coal only because it is more picturesque that way.

the equilibrium conditions have any explanatory value. For instance, the flow market that has to be cleared is not just one market; it is the sequence of markets for resource products from now until the date of exhaustion. It is, in other words, a sequence of futures markets, perhaps a long sequence. If the futures markets actually existed, we could perhaps accept the notion that their equilibrium configuration is stable; that might not be true, but it is at least the sort of working hypothesis we frequently accept as a way of getting on with business. But there clearly is not a full set of futures markets; natural-resource markets work with a combination of myopic flow transactions and rather more farsighted asset transactions. It is legitimate to ask whether observed resource prices are to be interpreted as approximations to equilibrium prices, or whether the equilibrium is so unstable that momentary prices are not only a bad indicator of equilibrium relationships, but also a bad guide to resource allocation.

That turns out not to be an easy question to answer. Flow considerations and stock considerations work in opposite directions. The flow markets by themselves could easily be unstable; but the asset markets provide a corrective force. Let me try to explain why.

The flow equilibrium condition is that the net price grow like compound interest at the prevailing rate. Suppose net prices are expected by producers to be rising too slowly. Then resource deposits are a bad way to hold wealth. Mine owners will try to pull out; and if they think only in flow terms, the way to get out of the resource business is to increase current production and convert ore into money. If current production increases, for this or any other reason, the current price must move down along the demand curve. So initially pessimistic price expectations on the part of producers have led to more pressure on the

current price. If expectations about future price changes are responsive to current events, the consequence can only be that pessimism is reinforced and deepened. The initial disequilibrium is worsened, not eliminated, by this chain of events. In other words, the market mechanism I have just described is unstable. Symmetrical reasoning leads to the conclusion that if prices are initially expected to be rising too fast, the withholding of supplies will lead to a speculative run-up of prices which is self-reinforcing. Depending on which way we start, initial disequilibrium is magnified, and production is tilted either toward excessive current dumping or toward speculative withholding of supply. (Still other assumptions are possible and may lead to qualitatively different results. For instance, one could imagine that expectations focus on the price level rather than its rate of change. There is much more work to be done on this question.)

Such things have happened in resource markets; but they do not seem always to be happening. I think that this story of instability in spot markets needs amendment; it is implausible because it leaves the asset market entirely out of account. The longer run prospect is not allowed to have any influence on current happenings. Suppose that producers do have some notion that the resource they own has a value anchored somewhere in the future, a value determined by technological and demand considerations, not by pure and simple speculation. Then if prices are now rising toward that rendezvous at too slow a rate, that is indeed evidence that owning resource deposits is bad business. But that will lead not to wholesale dumping of current production, but to capital losses on existing stocks. When existing stocks have been written down in value, the net price can rise toward its future rendezvous at more or less the right rate. As well as be-

ing destabilized by flow reactions, the market can be stabilized by capitalization reactions. In fact the two stories can be made to merge: the reduction in flow price coming from increased current production can be read as a signal and capitalized into losses on asset values, after which near-equilibrium is reestablished.

I think the correct conclusion to be drawn from this discussion is not that either of the stories is more likely to be true. It is more complex: that in tranquil conditions, resource markets are likely to track their equilibrium paths moderately well, or at least not likely to rush away from them. But resource markets may be rather vulnerable to surprises. They may respond to shocks about the volume of reserves, or about competition from new materials, or about the costs of competing technologies, or even about near-term political events, by drastic movements of current price and production. It may be quite a while before the transvaluation of values—I never thought I could quote Nietzsche in an economics paper—settles down under the control of sober future prospects. In between, it may be a cold winter.

So far, I have discussed the economic theory of exhaustible resources as a partial-equilibrium market theory. The interest rate that more or less controls the whole process was taken as given to the mining industry by the rest of the economy. So was the demand curve for the resource itself. And when the market price of the resource has ridden up the demand curve to the point where the quantity demanded falls to zero, the theory says that the resource in question will have been exhausted.

There is clearly a more cosmic aspect to the question than this; and I do not mean to suggest that it is unimportant, just because it is cosmic. In particular, there remains an important question about the

social interest in the pace of exploitation of the world's endowment of exhaustible natural resources. This aspect has been brought to a head recently, as everyone knows, by the various Doomsday forecasts that combine a positive finding that the world is already close to irreversible collapse from shortage of natural resources and other causes with the normative judgment that civilization is much too young to die. I do not intend to discuss those forecasts and judgments now—this convention already has one session devoted to just that—but I do want to talk about the economic issues of principle involved.

First, there is a proposition that will be second nature to everyone in this room. What I have called the fundamental principle of the economics of exhaustible resources is, among other things, a condition of competitive equilibrium in the sequence of futures markets for deliveries of the natural resource. This sequence extends out to infinity, even if the competitive equilibrium calls for the resource to be exhausted in finite time. Beyond the time of exhaustion there is also equilibrium: supply equals demand equals zero at a price simultaneously so high that demand is choked off and so low that it is worth no one's while to lose interest by holding some of the resource that long. Like any other competitive equilibrium with the right background assumptions, this one has some optimality properties. In particular, as Hotelling pointed out, the competitive equilibrium maximizes the sum of the discounted consumer-plus-producer surpluses from the natural resource, *provided* that society wishes to discount future consumer surpluses at the same rate that mine owners choose to discount their own future profits.

Hotelling was not so naive as to leap from this conclusion to the belief that *laissez-faire* would be an adequate policy for the resource industries. He pointed to

several ways in which the background assumptions might be expected to fail: the presence of externalities when several owners can exploit the same underground pool of gas or oil; the considerable uncertainty surrounding the process of exploration with the consequent likelihood of wasteful rushes to stake claims and exploit, and the creation of socially useless wind-fall profits; and, finally, the existence of large monopolistic or oligopolistic firms in the extractive industries.

There is an amusing sidelight here. It is not hard to show that, generally speaking, a monopolist will exhaust a mine more slowly than a competitive industry facing the same demand curve would do. (Hotelling did not explore this point in detail, though he clearly knew it. He did mention the possibility of an extreme case in which competition will exhaust a resource in finite time and a monopolist only asymptotically.) The amusing thing is that if a conservationist is someone who would like to see resources conserved *beyond* the pace that competition would adopt, then the monopolist is the conservationist's friend. No doubt they would both be surprised to know it.

Hotelling mentions, but rather poorly, the notion that market rates of interest might exceed the rate at which society would wish to discount future utilities or consumer surpluses. I think a modern economist would take that possibility more seriously. It is certainly a potentially important question, because the discount rate determines the whole tilt of the equilibrium production schedule. If it is true that the market rate of interest exceeds the social rate of time preference, then scarcity rents and market prices will rise faster than they "ought to" and production will have to fall correspondingly faster along the demand curve. Thus the resource will be exploited too fast and exhausted too soon.

The literature has several reasons for expecting that private discount rates might be systematically higher than the correct social rate of discount. They fall into two classes. The first class takes it more or less for granted that society ought to discount utility and consumption at the same rates as reflective individuals would discount their own future utility and consumption. This line of thought then goes on to suggest that there are reasons why this might not happen. One standard example is the fact that individuals can be expected to discount for the riskiness of the future, and some of the risks for which they will discount are not risks to society but merely the danger of transfers within the society. Since there is not a complete enough set of insurance markets to permit all these risks to be spread properly, market interest rates will be too high. Insecurity of tenure, as William Vickrey has pointed out, is a special form of uncertainty with particular relevance to natural resources.

A second standard example is the existence of various taxes on income from capital; since individuals care about the after-tax return on investment and society about the before-tax return, if investment is carried to the point where the after-tax yield is properly related to the rate of time preference, the before-tax profitability of investment will be too high. I have nothing to add to this discussion.

The other class of reasons for expecting that private discount rates are too high and will thus distort intertemporal decisions away from social optimality denies that private time preference is the right basis for intertemporal decisions. Frank Ramsey, for instance, argued that it was ethically indefensible for society to discount future utilities. Individuals might do so, either because they lack imagination (Böhm-Bawerk's "defective telescopic faculty") or because they are all too con-

scious that life is short. In social decision-making, however, there is no excuse for treating generations unequally, and the time-horizon is, or should be, very long. In solemn conclave assembled, so to speak, we ought to act as if the social rate of time preference were zero (though we would simultaneously discount future *consumption* if we expect the future to be richer than the present). I confess I find that reasoning persuasive, and it provides another reason for expecting that the market will exhaust resources too fast.

This point need not be divorced so completely from individual time preference. If the whole infinite sequence of futures markets for resource products could actually take place and find equilibrium, I might be inclined to accept the result (though I would like to know who decides the initial endowments within and between generations). But of course they cannot take place. There is no way to collect bids and offers from everyone who will ever live. In the markets that actually do take place, future generations are represented only by us, their eventual ancestors. Now generations overlap, so that I worry about my children, and they about theirs, and so on. But it does seem fundamentally implausible that there should be anything *ex post* right about the weight that is actually given to the welfare of those who will not live for another thousand years. We have actually done quite well at the hands of *our* ancestors. Given now poor they were and how rich we are, they might properly have saved less and consumed more. No doubt they never expected the rise in income per head that has made us so much richer than they ever dreamed was possible. But that only reinforces the point that the future may be too important to be left to the accident of mistaken expectations and the ups and downs of the Protestant ethic.

Several writers have studied directly the

problem of defining and characterizing a socially-optimal path for the exploitation of a given pool of exhaustible resources. The idea is familiar enough: instead of worrying about market responses, one imagines an idealized planned economy, constrained only by its initial endowment, the size of the labor force, the available technology, and the laws of arithmetic. The planning board then has to find the best feasible development for the economy. To do so, it needs a precise criterion for comparing different paths, and that is where the social rate of time preference plays a role.

It turns out that the choice of a rate of time preference is even more critical in this situation than it is in the older literature on optimal capital accumulation without any exhaustible resources. In that theory, the criterion usually adopted is the maximization of a discounted sum of one-period social welfare indicators, depending on consumption per head, and summed over all time from now to the infinite future. The typical result, depending somewhat on the particular assumptions made, is that consumption per head rises through time to a constant plateau defined by the "modified Golden Rule." In that ultimate steady state, consumption per head is lower the higher is the social rate of discount; and, correspondingly, the path to the steady state is characterized by less saving and more interim consumption, the higher the social rate of discount. That is as it should be: the main beneficiaries of a high level of ultimate steady-state consumption are the inhabitants of the distant future, and so, if the planning board discounts the future very strongly, it will choose a path that favors the near future over the distant future.

When one adds exhaustible resources to the picture, the social rate of time preference can play a similar, but even more critical, role. As a paper by Geoffrey Heal

and Partha Dasgupta and one of my own show, it is possible that the optimal path with a positive discount rate should lead to consumption per head going asymptotically to zero, whereas a zero discount rate leads to perpetually rising consumption per head. In other words, even when the technology and the resource base could permit a plateau level of consumption per head, or even a rising standard of living, positive social time preference might in effect lead society to prefer eventual extinction, given the drag exercised by exhaustible resources. Of course, it is part of the point that it is the planning board in the present that plans for future extinction: nobody has asked the about-to-become-defunct last generation whether it approved of weighting its satisfactions less than those of its ancestors.

Good theory is usually trying to tell you something, even if it is not the literal truth. In this context, it is not hard to interpret the general tenor of the theoretical indications. We know in general that even well-functioning competitive markets may fail to allocate resources properly over time. The reason, I have suggested, is because, in the nature of the case, the future brings no endowment of its own to whatever markets actually exist. The intergenerational distribution of income or welfare depends on the provision that each generation makes for its successors. The choice of a social discount rate is, in effect, a policy decision about that intergenerational distribution. What happens in the planning parable depends very much—perhaps dramatically—on that choice; and one's evaluation of what happens in the market parable depends very much on whether private choices are made with a discount rate much larger than the one a deliberate policy decision would select. The pure theory of exhaustible resources is trying to tell us that, if exhaustible resources really matter, then the balance

between present and future is more delicate than we are accustomed to think; and then the choice of a discount rate can be pretty important and one ought not to be too casual about it.

In my own work on this question, I have sometimes used a rather special criterion that embodies sharp assumptions about intergenerational equity: I have imposed the requirement that consumption per head be constant through time, so that no generation is favored over any other, and asked for the largest steady consumption per head that can be maintained forever, given all the constraints including the finiteness of resources. This criterion, like any other, has its pluses and its minuses and I am not committed to it by any means. Like the standard criterion—the discounted sum of one-period utilities—this one will always pick out an *efficient* path, so one at least gets the efficiency conditions out of the analysis. The highest-constant-consumption criterion also has the advantage of highlighting the crucial importance of certain technological assumptions.

It is clear without any technical apparatus that the seriousness of the resource-exhaustion problem must depend in an important way on two aspects of the technology: first, the likelihood of technical progress, especially natural-resource-saving technical progress, and, second, the ease with which other factors of production, especially labor and reproducible capital, can be substituted for exhaustible resources in production.

My own practice, in working on this problem, has been to treat as the central case (though not the only case) the assumption of zero technological progress. This is not because I think resource-saving inventions are unlikely or that their capacity to save resources is fundamentally limited. Quite the contrary—if the future is anything like the past, there will be pro-

longed and substantial reductions in natural-resource requirements per unit of real output. It is true, as pessimists say, that it is just an assumption and one cannot be sure; but to assume the contrary is also an assumption, and a much less plausible one. I think there is virtue in analyzing the zero-technical-progress case because it is easy to see how technical progress can relieve and perhaps eliminate the drag on economic welfare exercised by natural-resource scarcity. The more important task for theory is to try to understand what happens or can happen in the opposite case.

As you would expect, the degree of substitutability is also a key factor. If it is very easy to substitute other factors for natural resources, then there is in principle no "problem." The world can, in effect, get along without natural resources, so exhaustion is just an event, not a catastrophe. Nordhaus's notion of a "backstop technology" is just a dramatic way of putting this case; at some finite cost, production can be freed of dependence on exhaustible resources altogether.

If, on the other hand, real output per unit of resources is effectively bounded—cannot exceed some upper limit of productivity which is in turn not too far from where we are now—then catastrophe is unavoidable. In-between there is a wide range of cases in which the problem is real, interesting, and not foreclosed. Fortunately, what little evidence there is suggests that there is quite a lot of substitutability between exhaustible resources and renewable or reproducible resources, though this is an empirical question that could absorb a lot more work than it has had so far.

Perhaps the most dramatic way to illustrate the importance of substitutability, and its connection with Doomsday, is in terms of the permanent sustainability of a constant level of consumption. In the

simplest, most aggregative, model of a resource-using economy one can prove something like the following: if the elasticity of substitution between exhaustible resources and other inputs is unity or bigger, and if the elasticity of output with respect to reproducible capital exceeds the elasticity of output with respect to natural resources, then a constant population can maintain a positive constant level of consumption per head forever. This permanently maintainable standard of living is an increasing, concave, and unbounded function of the initial stock of capital. So the drag of a given resource pool can be overcome *to any extent* if only the initial stock of capital is large enough. On the other hand, if the elasticity of substitution between natural resources and other inputs is less than one, or if the elasticity of output with respect to resources exceeds the elasticity of output with respect to reproducible capital, then the largest constant level of consumption sustainable forever with constant population is—zero. We know much too little about which side of that boundary the world is on—technological progress aside—but at least the few entrails that have been read seem favorable.²

Perhaps I should mention that when I say "forever" in this connection, I mean "for a very long time." The mathematical reasoning does deal with infinite histories, but actually life in the solar system will only last for a finite time, though a very long finite time, much longer than this lecture, for instance. That is why I think it takes economics as well as the entropy law to answer our question.

I began this lecture by talking of the conditions for competitive equilibrium in the market for natural resources. Now I have been talking of centralized planning

² See pp. 60–70 in William D. Nordhaus and James Tobin.

optima. As you would expect, it turns out that under the standard assumptions, the Hotelling rule, the fundamental principle of natural-resource economics, is a necessary condition for efficiency and therefore for social optimality. So there is at least a prayer that a market-guided system might manage fairly well. But more than the Hotelling condition is needed.

I have already mentioned one of the extra requirements for the intertemporal optimality of market allocations: it is that the market discount future profits at the same rate as the society would wish to discount the welfare of future inhabitants of the planet. This condition is often given as an argument for public intervention in resource allocation because—as I have also mentioned—there are reasons to expect market interest rates to exceed the social rate of time preference, or at least what philosophers like us think it ought to be. If the analysis is right, then the market will tend to consume exhaustible resources too fast, and corrective public intervention should be aimed at slowing down and stretching out the exploitation of the resource pool. There are several ways that could be done, in principle, through conservation subsidies or a system of graduated severance taxes, falling through time.

Realistically speaking, however, when we say "public intervention" we mean rough and ready political action. An only moderately cynical observer will see a problem here: it is far from clear that the political process can be relied on to be more future-oriented than your average corporation. The conventional pay-out period for business is of the same order of magnitude as the time to the next election, and transferring a given individual from the industrial to the government bureaucracy does not transform him into a guardian of the far future's interests. I have no ready solution to this problem. At a minimum, it suggests that one ought to be as

suspicious of uncritical centralization as of uncritical free-marketeering. Maybe the safest course is to favor specific policies—like graduated severance taxes—rather than blanket institutional solutions.

There is another, more subtle, extra requirement for the optimality of the competitive market solution to the natural-resource problem. Many patterns of exploitation of the exhaustible-resource pool obey Hotelling's fundamental principle myopically, from moment to moment, but are wrong from a very long-run point of view. Such mistaken paths may even stay very near the right path for a long time, but eventually they veer off and become bizarre in one way or another. If a market-guided system is to perform well over the long haul, it must be more than myopic. Someone—it could be the Department of the Interior, or the mining companies, or their major customers, or speculators—must always be taking the long view. They must somehow notice in advance that the resource economy is moving along a path that is bound to end in disequilibrium of some extreme kind. If they do notice it, and take defensive actions, they will help steer the economy from the wrong path toward the right one.⁴ Usually the "wrong" path is one that leads to exhaustion at a date either too late or too soon; anyone who perceives this will be motivated to arbitrage between present and future in ways that will push the current price toward the "right" path.⁵

It is interesting that this need for some-

⁴ This sort of process has been studied in a different context by Frank Hahn and by Karl Shell and Joseph Stiglitz.

⁵ For example, suppose the current price is too low, in the sense that, if it rises according to the current principle, the demand path will be enough to exhaust the resource before the price has risen high enough to choke demand to zero. A clever speculator would see that there will be money to be made just after the date of exhaustion, because anyone with a bit of the resource to sell could make a discrete jump in the price and still find buyers. Such a speculator would wish to buy now

one to take the long view emerged also when the question at hand was the potential instability of the market for natural resources if it concentrates too heavily on spot or flow decisions, and not enough on future or stock decisions. In that context too, a reasonably accurate view of the long-term prospects turns out to be a useful, maybe indispensable, thing for the resource market to have.

This lecture has been—as Kenneth Burke once said about the novel—words, all words. Nevertheless, it has been a discourse on economic theory, not on current policy. If some of you have been daydreaming about oil and the coming winter, I assure you that I have been thinking about shadow prices and transversality conditions at infinity. If I turn briefly to policy at the end, it is not with concrete current problems in mind. After all, nothing I have been able to say takes account of the international oil cartel, the political and economic ambitions of the Middle Eastern potentates, the speeds of adjustment to surprises in the supply of oil, or the doings of our own friendly local oligopolists. The only remarks I feel entitled to make are about the long-run pursuit of a general policy toward exhaustible resources.

Many discussions of economic policy—macroeconomics aside—boil down to a tension between market allocation and public intervention. Marketeers keep thinking about the doughnut of allocative efficiency and informational economy and *dirigistes* are impressed with the size of the hole containing externalities, imperfections, and distributional issues. So it is with exhaustible resources. One is impressed with what a system of ideal mar-

kets, including futures markets, can accomplish in this complicated situation; and one can hardly miss seeing that our actual oligopolistic, politically involved, pollution-producing industry is not exactly what the textbook ordered. I have nothing new to add to all that. The unusual factor that the theory of exhaustible resources brings to the fore is the importance of the long view, and the value of reasonable information about reserves, technology, and demand in the fairly far future.

This being so, one is led to wonder whether public policy can contribute to stability and efficiency along those lines. One possibility is the encouragement of organized futures trading in natural resource products. To be useful, futures contracts would have to be much longer-term than is usual in the futures markets that now exist, mostly for agricultural products. I simply do not know enough to have an opinion about the feasibility of large scale futures trading, or about the ultimate contribution that such a reform would make to the stability and efficiency of the market for resource products. But in principle it would seem to be a good idea.

The same considerations suggest that the market for exhaustible resources might be one of the places in the economy where some sort of organized indicative planning could play a constructive role. This is not an endorsement of centralized decision-making, which is likely to have imperfections and externalities of its own. Indeed it might be enough to have the government engaged in a continuous program of information-gathering and dissemination covering trends in technology, reserves and demand. One could at least hope to have professional standards govern such an exercise. I take it that the underlying logic of indicative planning is that some comparison and coordination of the main participants in the market, including the

and hold for sale then. But that action would tend to raise the current price (and, by the fundamental principle, the whole price path) and reduce demand, so that the life of the resource would be prolonged. The speculation is thus corrective.

government, could eliminate major errors and resolve much uncertainty. In the case of exhaustible resources, it could have the additional purpose of generating a set of consistent expectations about the distant future. In this effort, the pooling of information and intentions from both sides of the market could be useful, with the effect of inducing behavior that will steer the economy away from ultimately inferior paths. It is also likely, as Adam Smith would have warned, that a certain amount of conspiracy against the public interest might occur in such sessions, so perhaps they ought to be recorded and the tapes turned over to Judge Sirica, who will know what to do with them.

REFERENCES

- P. Dasgupta and G. Heal, "The Optimal Depletion of Exhaustible Resources," *Rev. Econ. Stud.*, forthcoming, 1974.
- F. H. Hahn, "Equilibrium Dynamics with Heterogeneous Capital Goods," *Quart. J. Econ.*, Nov. 1966, 80, 633-646.
- H. Hotelling, "The Economics of Exhaustible Resources," *J. Polit. Econ.*, April 1931, 39, 137-175.
- W. D. Nordhaus, "The Allocation of Energy Resources," *Brookings Papers on Econ. & Activ.*, forthcoming.
- and J. Tobin, "Is Economic Growth Obsolete?" in National Bureau of Economic Research, *Economic Growth*, 50th Anniversary Colloq. V, New York 1972.
- K. Shell and J. E. Stiglitz, "The Allocation of Investment in a Dynamic Economy," *Quart. J. Econ.*, Nov. 1967, 81.
- R. M. Solow, "Intergenerational Equity and Exhaustible Resources," *Rev. Econ. Stud.*, forthcoming, 1974.
- J. E. Stiglitz, "Growth with Exhaustible Natural Resources," *Rev. Econ. Stud.*, forthcoming, 1974.
- M. Weinstein and R. Zeckhauser, "Use Patterns for Depletable and Recyclable Resources," *Rev. Econ. Stud.*, forthcoming, 1974.