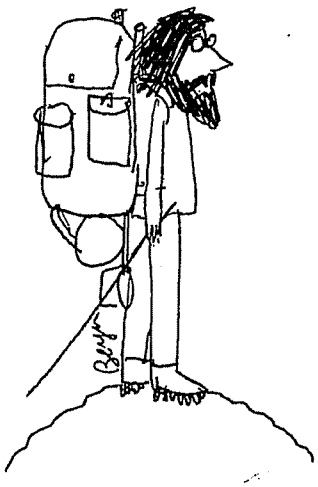

ROBERT M. SOLOW

Is the End of the World at Hand?

To grow or not to grow: that never was the question until recently. But a computer model developed by MIT's Jay Forrester projects imminent collapse for the world economy unless we call a halt to growth now. Professor Solow has some doubts which he expressed at a Symposium on the Limits to Growth held at Lehigh University.

I was having a hard time figuring out how to begin when I came across an excerpt from an interview with my MIT colleague Professor Jay Forrester, who is either the Christopher Columbus or the Dr. Strangelove of this business, depending on how you look at it. Forrester said he would like to see about 100 individuals, the most gifted and best qualified in the world, brought together in a team to make a psychosocial analysis of the problem of world equilibrium. He thought it would take about ten years. When he was asked to define the composition of his problem-solving group, Forrester said: "Above all it shouldn't be mostly made up of professors. One would include people who had been successful in their personal careers, whether in politics, business, or anywhere else. We should also need radical philosophers, but we should take care to keep out representatives of the social sciences. Such people always want to go to the bottom of a particular problem. What we want to look at are the problems caused by interactions."

I don't know what you call people who believe they can be wrong about everything in particular, but expect to be lucky enough somehow to get it right on the interactions. They may be descendants of the famous merchant Lapidus, who said he lost money on every item he sold, but made it up on the volume. Well, I suppose that as an economist

ROBERT M. SOLOW is Professor of Economics at MIT. His paper, along with others presented at Lehigh University, will appear in *The Economic Growth Controversy*, to be published this spring by International Arts & Sciences Press, Inc.

I am a representative of the social sciences; and I'm prepared to play out the role by talking about first principles and trying to say what the Growth vs. No-Growth business is really all about. This is going to involve me in the old academic ploy of saying over and over again what I'm not talking about before I ever actually say what I think I am talking about. But I'm afraid that some of those boring distinctions are part of the price you have to pay for getting it right.

First of all, there are (at least) two separate questions you can ask about the prospects for economic growth. You can ask: Is growth desirable? Or you can ask: Is growth possible? I suppose that if continued economic growth is not possible, it hardly matters whether or not it's desirable. But if it is possible, it's presumably not inevitable, so we can discuss whether we should want it. But they are separate questions, and an answer to one of them is not necessarily an answer to the other. My main business is with the question about the possibility of continued growth; I want to discuss the validity of the negative answer given by the "Doomsday Models" associated with the names of Forrester and Meadows (and MIT!) and, to a lesser extent, with the group of English scientists who published a manifesto called "Blueprint for Survival." The main concern of Dr. E. J. Mishan [whose article will appear in a later issue], on the other hand, was with the desirability of continued economic growth (and, at least by implication, with the desirability of past economic growth). If I spend a few minutes poaching on his territory, it is mainly because that seems like a good way to get some concepts straight, but also just to keep a discussion going.

Sorting out the issues

Arguments about the desirability of economic growth often turn quickly into arguments about the "quality" of modern life. One gets the notion that you favor growth if you are the sort of person whose idea of heaven is to drive at 90 miles an hour down a six-lane highway reading billboards, in order to pollute the air over some crowded lake with the exhaust from twin 100-horsepower outboards, and whose idea of food is Cocoa Krispies. On the other hand, to be against economic growth is to be a granola-eating, backpacking, transcendental-meditating canoe freak. That may even be a true statistical association, but I will argue that there is no necessary or logical connection between your answer to the growth question and your answer to the quality-of-life question. Suppose there were no issue about economic growth; suppose it were impossible; suppose each man or each woman were equipped to have only two children (one bomb under each wing); suppose we were stuck with the technology we have now and had no concept of invention, or even of increased mechanization through capital investment. We could still argue about the relative merits of cutting timber for building houses or leaving it stand to be enjoyed as forest. Some people would still be willing to breathe carbon monoxide in big cities in return for the excitement of urban life, while others would prefer cleaner air and fewer TV channels. Macy's would still not tell Gimbel's. Admen would still try to tell you that all those beautiful women are actually just looking for somebody who smokes Winchesters, thus managing to insult both men and women at once. Some people would still bring transistor radios to the beach. All or nearly all of the arguments about the quality of life would be just as valid if the question of growth never arose.

I won't go so far as to say there is no connection. In particular, one can argue that if population density were low enough, people would interfere much less with each other, and everyone could find a part of the world and style of civilization that suited him. Then differences of opinion about the quality of life wouldn't matter so much. Even if I grant the truth of that observation, it is still the case that, from here on out, questions about the quality of life are separable from questions about the desirability of growth. If growth stopped, there would be just about as much to complain about; and, as I shall argue later on, one can imagine continued growth that is directed against pollution, against congestion, against sliced white bread.

I suppose it is only fair to admit that if you get very enthusiastic about economic growth you are likely to be attracted to easily quantifiable and measurable things as objects of study, to point at with pride or to view with alarm. You are likely to pay less attention to important, intangible aspects of the standard of living. Although you can't know whether people are happier than they used to be, you can at least determine that they drink more

orange juice or take more aspirin. But that's mere weakness of imagination and has nothing to do in principle with the desirability of economic growth, let alone with its possibility.

There is another practical argument that is often made; and although it is important, it sometimes serves as a way of avoiding coming to grips with the real issues. This argument says that economic growth, increasing output per person, is the only way we are likely to achieve a more equitable distribution of income in society. There is a lot of home truth in that. It is inevitably less likely that a middle-class electorate will vote to redistribute part of its own income to the poor than that it will be willing to allocate a slightly larger

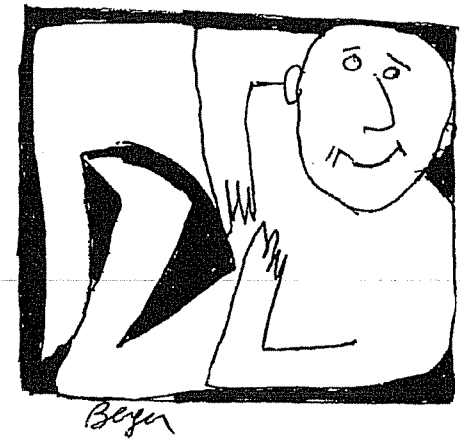

share of a growing total. Even more pessimistically, I might suggest that even a given relative distribution of income, supposing it cannot be made more nearly equal, for political or other reasons, is less unattractive if the absolute standard of living at the bottom is fairly high than it is if the absolute standard at the bottom is very low. From this point of view, even if economic growth doesn't lead to more equity in distribution, it makes the inequity we've got more tolerable. I think it is one of the lessons of history as recent as the McGovern campaign that this is a realistic statement of the prospects.

It is even clearer if one looks, not at the distribution of income within a rich country like the U.S., but at the distribution of income between the developed countries of the world and the undeveloped ones. The rich Western nations have never been able to agree on the principle of allocating as much as one percent of their GNP to aid undeveloped countries. They are unlikely to be willing to share their wealth on any substantial scale with the poor countries. Even if they were, there are so many more poor people in the world that an equally shared income would be quite low. The *only* prospect of a decent life for Asia, Africa, and Latin America is in more total output.

But I point this out only to warn you that it is not the heart of the question. I think that those who oppose continued growth should in honesty face up to the implications of their position for distributional equity and the prospects of the world's poor. I think those who favor continued growth on the grounds that only thus can we achieve some real equality ought to be serious about that. If economic growth with equality is a good thing, it doesn't follow that economic growth with a lot of pious talk about equality is a good thing. In principle, we can have growth with or without equity; and we can have stagnation with or without equity. An argument about first principles should keep those things separate.

What has posterity done for us?

Well, then, what *is* the problem of economic growth all about? (I'm giving a definition now, not stating a fact, so all I can say is that I think this way of looking at it contributes to clarity of thought.) Whenever there is a question about what to *do*, the desirability of economic growth turns on the claims of the future against the claims of the present. The pro-growth-man is someone who is prepared to sacrifice something useful and desirable right now so that people should be better off in the future; the anti-growth-man is someone who thinks that is unnecessary or undesirable. The nature of the sacrifice of present enjoyment for future enjoyment can be almost anything. The classic example is investment: We can use our labor and our resources to build very durable things like roads or subways or factories or blast furnaces or dams that will be used for a long time by people who were not even born when those things were created, and so will certainly have contributed nothing to their construction. That labor and those resources can just as well be used to produce shorter-run pleasures for us now.

Such a sacrifice of current consumption on behalf of the future may not strike you as much of a

sacrifice. But that's because you live in a country that is already rich; if you had lived in Stalin's Russia, that need to sacrifice would be one of the reasons you would have been given to explain why you had to live without comfort and pleasures while the Ministry of Heavy Industry got all the play. If you lived in an underdeveloped country now you would face the same problem: What shall you do with the foreign currency earned by sales of cocoa or copper or crude oil—spend it on imports of consumer goods for those alive and working now, or spend it on imports of machinery to start building an industry that may help to raise the standard of living in 30 years' time?

There are other ways in which the same choice can be made, including, for instance, the direction of intellectual resources to the invention of things (like the generation of electricity from nuclear fusion) that will benefit future generations. Paradoxically, one of the ways in which the present can do something for the future is to conserve natural resources. If we get along with less lumber now so that there will be more forests standing for our grandchildren, or if we limit the present consumption of oil or zinc so that there will be some left for the twenty-first century, or if we worry about siltation behind dams that would otherwise be fun for fishermen and water-skiers, in all those cases we are promoting economic growth. I call that paradoxical because I think most people identify the conservation freak with the anti-growth party whereas, in this view of the matter, the conservationist is trading present satisfaction for future satisfaction, that is, he is promoting economic growth. I think the confusion comes from mixing up the quality-of-life problem with the growth problem. But it is nonetheless a confusion.

Why should we be concerned with the welfare of posterity, given the indubitable fact that posterity has never done a thing for us? I am not anthropologist enough to know how rare or common it is that our culture should teach us to care not only about our children but about their children, and their children. I suppose there are good Darwinian reasons why cultures without any future-orientation should fail to survive very long in the course of history. (But remember that they had a merry time of it while they lasted!) Moreover, we now enjoy the investments made by our ancestors, so there is a kind of equity in passing it on. Also, unless something terrible happens, there will be a lot more future than there has been past; and, for better or worse—probably worse—there will be more people at each future instant than there are now or have been. So all in all, the future will involve many more man-years of life than the present or the past, and a kind of intergenerational democracy suggests that all those man-years-to-be deserve some consideration out of sheer numbers.

On the other hand, *if* continued economic growth is possible—which is the question I'm coming to—then it is very likely that posterity will be richer than we are even if we make no special efforts on its behalf. If history offers any guide, then, in the developed part of the world at least, the accumulation of technological knowledge will probably make our great-grandchildren better off than we are, even if we make no great effort in that direction. Leaving aside the possibility of greater equality—I have already discussed that—there is hardly a crying need for posterity to be on average very much richer than we are. Why should we poor folk make any sacrifices for those who will in any case live in luxury in the future? Of course, if the end of the world is at hand, if continued economic growth is *not* possible, then we ought to care more about posterity, because they won't be so well off. Paradoxically, if continued growth is not possible, or less possible, then we probably ought to do more to promote it. Actually, there's no paradox in that, as every student of economics will realize, because it is a way of saying that the marginal return on investment is high.

Overshoot, collapse, doom

There is, as you know, a school of thought that claims that continued economic growth is in fact not possible anymore, or at least not for very long. This judgment has been expressed more or less casually by several observers in recent years. What distinguishes the "Doomsday Models" from their predecessors is that they claim to much more than a casual judgment: they deduce their beliefs about future prospects from mathematical models or systems analysis. They don't merely say that the end of the world is at hand—they can show you computer output that says the same thing.

Characteristically, the Doomsday Models do more than just say that continued economic growth

is impossible. They tell us why: in brief, because (*a*) the earth's natural resources will soon be used up; (*b*) increased industrial production will soon strangle us in pollution; and (*c*) increasing population will eventually outrun the world's capacity to grow food, so that famine must eventually result. And, finally, the models tell us one more thing: the world will end with a bang, not a whimper. The natural evolution of the world economy is not at all toward some kind of smooth approach to its natural limits, wherever they are. Instead, it is inevitable—unless we make drastic changes in the way we live and organize ourselves—that the world will overshoot any level of population and production it can possibly sustain and will then collapse, probably by the middle of the next century.

I would like to say why I think that the Doomsday Models are bad science and therefore bad guides to public policy. I hope nobody will conclude that I believe the problems of population control, environmental degradation, and resource exhaustion to be unimportant, or that I am one of those people who believe that an adequate response to such problems is a vague confidence that some technological solution will turn up. On the contrary, it is precisely because these are important problems that public policy had better be based on sound and careful analysis. I want to explain some of my reasons for believing that the global models don't provide even the beginnings of a foundation of that kind.

The first thing to realize is that the characteristic conclusion of the Doomsday Models is very near the surface. It is, in fact, more nearly an assumption than a conclusion, in the sense that the chain of logic from the assumptions to the conclusion is very short and rather obvious.

The basic assumption is that stocks of things like the world's natural resources and the waste-disposal capacity of the environment are finite, that the world economy tends to consume the stock at an increasing rate (through the mining of minerals and the production of goods), and that there are no built-in mechanisms by which approaching exhaustion tends to turn off consumption gradually and in advance. You hardly need a giant computer to tell you that a system with those behavior rules is going to bounce off its ceiling and collapse to a low level. Then, in case anyone is inclined to relax into the optimistic belief that maybe things

aren't that bad, we are told: Imagine that the stock of natural resources were actually twice as big as the best current evidence suggests, or imagine that the annual amount of pollution could be halved all at once and then set to growing again. All that would happen is that the date of collapse would be postponed by T years, where T is not a large number. But once you grasp the quite simple essence of the models, this should come as no surprise. It is important to realize where these powerful conclusions come from, because, if you ask yourself "Why didn't I realize earlier that the end of the world was at hand?" the answer is not that you weren't clever enough to figure it out for yourself. The answer is that the imminent end of the world is an immediate deduction from certain assumptions, and one must really ask if the assumptions are any good.

It is a commonplace that if you calculate the annual output of any production process, large or small, and divide it by the annual employment of labor, you get a ratio that is called the productivity of labor. At the most aggregative level, for example, we can say that the GNP in 1971 was $1,050 billion and that about 82 million people were employed in producing it, so that GNP per worker or the productivity of a year of labor was about $12,800. Symmetrically, though the usage is less common, one could just as well calculate the GNP per unit of some particular natural resource and call that the productivity of coal, or GNP per pound of vanadium. We usually think of the productivity of labor as rising more or less exponentially, say at 2 or 3 percent a year, because that is the way it has in fact behaved over the past century or so since the statistics began to be collected. The rate of increase in the productivity of labor is not a constant of nature. Sometimes it is faster, sometimes slower. For example, we know that labor productivity must have increased more slowly a long time ago, because if we extrapolate backward at 2 percent a year, we come to a much lower labor productivity in 1492 than can possibly have been the case. And the productivity of labor has risen faster in the past 25 years than in the 50 years before that. It also varies from place to place, being faster in Japan and Germany and slower in Great Britain, for reasons that are not at all certain. But it rises, and we expect it to keep rising.

Now, how about the productivity of natural resources? All the Doomsday Models will allow is a one-time hypothetical increase in the world supply of natural resources, which is the equivalent of a one-time increase in the productivity of natural resources. Why shouldn't the productivity of most natural resources rise more or less steadily through time, like the productivity of labor?

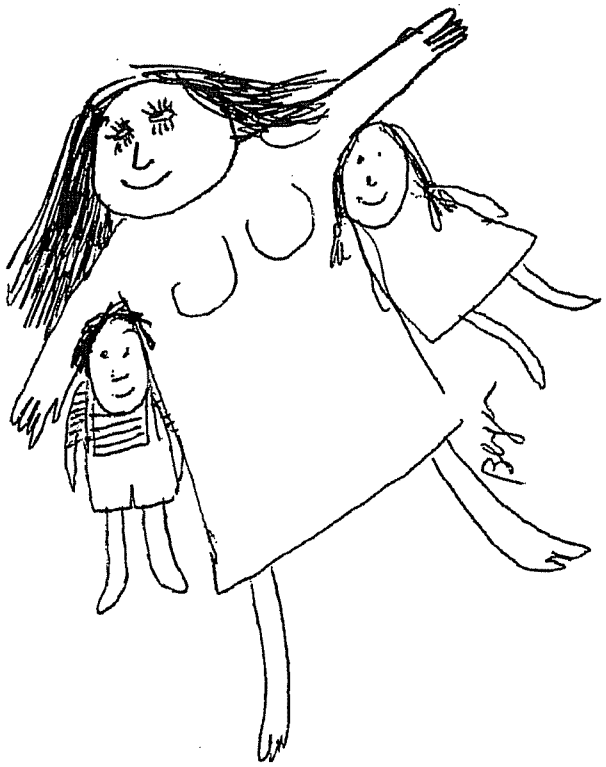

Of course it does for some resources, but not for others. Real GNP roughly doubled between 1950 and 1970. But the consumption of primary and scrap iron increased by about 20 percent, so the productivity of iron, GNP per ton of iron, increased by about 2.5 percent a year on the average during those 20 years. The U.S. consumption of manganese rose by 30 percent in the same period, so the productivity of manganese went up by some 70 percent in 20 years, a bit under 2.25 percent a year. Aggregate consumption of nickel just about doubled, like GNP, so the productivity of nickel didn't change. U.S. consumption of copper, both primary and secondary, went up by a third between 1951 and 1970, so GNP per pound of copper rose at 2 percent a year on the average. The story on lead and zinc is very similar, so their productivity increased at some 2 percent a year. The productivity of bituminous coal rose at 3 percent a year.

Naturally, there are important exceptions, and unimportant exceptions. GNP per barrel of oil was about the same in 1970 as in 1951: no productivity increase there. The consumption of natural gas tripled in the same period, so GNP per cubic foot of natural gas fell at about 2.5 percent a year. Our industrial demand for aluminum quadrupled in two decades, so the productivity of aluminum fell at a good 3.5 percent a year. And industrial demand for columbium was multiplied by a factor of 25: in 1951 we managed $2.25 million of GNP (in 1967 prices) per pound of columbium, whereas in 1970 we were down to $170 thousand of GNP per pound of columbium. On the other hand, it is a little hard to imagine civilization toppling because of a shortage of columbium.

Obviously many factors combine to govern the course of the productivity of any given mineral over time. When a rare natural resource is first available, it acquires new uses with a rush; and consumption goes up much faster than GNP. That's the columbium story, no doubt, and, to a lesser extent, the vanadium story. But once the novelty has worn off, the productivity of a resource tends to rise as better or worse substitutes for it appear, as new commodities replace old ones, and as manufacturing processes improve. One of the reasons the productivity of copper rises is because that of aluminum falls, as aluminum replaces copper in many uses. The same is true of coal and oil. A resource, like petroleum, which is versatile because of its role as a source of energy, is an interesting special case. It is hardly any wonder that the productivity of petroleum has stagnated, because the consumption of energy—both as electricity for domestic and industrial use and in the automobile—has recently increased even faster than GNP. But no one can doubt that we will run out of oil, that coal and nuclear fission will replace oil as the major sources of energy. It is already becoming probable that the high-value use of oil will soon be as feed stock for the petrochemical industries, rather than as a source of energy. Sooner or later, the productivity of oil will rise out of sight, because the production and consumption of oil will eventually dwindle toward zero, but real GNP will not.

So there really is no reason why we should not think of the productivity of natural resources as increasing more or less exponentially over time.

But then overshoot and collapse are no longer the inevitable trajectory of the world system, and the typical assumption-conclusion of the Doomsday Models falls by the wayside. We are in a different sort of ball game. The system might still burn itself out and collapse in finite time, but one cannot say with any honesty that it must. It all depends on the particular, detailed facts of modern economic life as well as on the economic policies we and the rest of the world pursue. I don't want to argue for any particular counterstory; all I want to say now is that the overshoot-collapse pattern is built into the models very near the surface, by assumption, and by implausible assumption at that.

Scarcity—and high prices

There is at least one reason for believing that the Doomsday story is almost certainly wrong. The most glaring defect of the Forrester-Meadows models is the absence of any sort of functioning price system. I am no believer that the market is always right, and I am certainly no advocate of laissez-faire where the environment is concerned. But the price system is, after all, the main social institution evolved by capitalist economies (and, to an increasing extent, socialist economies too) for registering and reacting to relative scarcity. There are several ways that the working of the price system will push our society into faster and more systematic increases in the productivity of natural resources.

First of all, let me go back to the analogy between natural resources and labor. We are not surprised to learn that industry quite consciously tries to make inventions that save labor, i.e., permit the same product to be made with fewer man-hours of work. After all, on the average, labor costs amount to almost three-fourths of all costs in our economy. An invention that reduces labor requirements per unit of GNP by 1 percent reduces all costs by about 0.75 percent. Natural resource costs are a much smaller proportion of total GNP, something nearer 5 percent. So industry and engineering have a much stronger motive to reduce labor requirements by 1 percent than to reduce resource requirements by 1 percent, assuming—which may or not be true—that it is about as hard to do one

as to do the other. But then, as the earth's supply of particular natural resources nears exhaustion, and as natural resources become more and more valuable, the motive to economize those natural resources should become as strong as the motive to economize labor. The productivity of resources should rise faster than now—it is hard to imagine otherwise.

There are other ways in which the market mechanism can be expected to push us all to economize on natural resources as they become scarcer. Higher and rising prices of exhaustible resources lead competing producers to substitute other materials that are more plentiful and therefore cheaper. To the extent that it is impossible to design around or find substitutes for expensive natural resources, the prices of commodities that contain a lot of them will rise relative to the prices of other goods and services that don't use up a lot of resources. Consumers will be driven to buy fewer resource-intensive goods and more of other things. All these effects work automatically to increase the productivity of natural resources, i.e., to reduce resource requirements per unit of GNP.

As I mentioned a moment ago, this is not an argument for laissez-faire. We may feel that the private decisions of buyers and sellers give inadequate representation to future generations. Or we may feel that private interests are in conflict with a distinct public interest—strip-mining of coal is an obvious case in point, and there are many others as soon as we begin to think about environmental effects. Private market responses may be too uncoordinated, too slow, based on insufficient and faulty information. In every case there will be actions that public agencies can take and should take; and it will be a major political struggle to see that they are taken. But I don't see how one can have the slightest confidence in the predictions of models that seem to make no room for the operation of everyday market forces. If the forecasts are wrong, then so are the policy implications, to the extent that there are any realistic policy implications.

Every analysis of resource scarcity has to come to terms with the fact that the prices of natural resources and resource products have not shown any tendency to rise over the past half-century, relative to the prices of other things. This must mean that there have so far been adequate offsets to any progressive impoverishment of deposits—like improvements in the technology of extraction, savings in end uses, or the availability of cheaper substitutes. The situation could, of course, change; and very likely some day it will. If the experienced and expert participants in the market now believed that resource prices would be sharply higher at some foreseeable time, prices would *already* be rising, as I will try to explain in a moment. The historical steadiness of resource prices suggests that buyers and sellers in the market have not been acting as if they foresaw exhaustion in the absence of substitutes, and therefore sharply higher future prices. They may turn out to be wrong; but the Doomsday Models give us absolutely no reason to expect that—in fact, they claim to get whatever meager empirical basis they have from such experts.

Why is it true that if the market saw higher prices in the future, prices would already be rising? It is a rather technical point, but I want to explain it because, in a way, it summarizes the important thing about natural resources: conserving a mineral deposit is just as much of an investment as building a factory, and it has to be analyzed that way. Any owner of a mineral deposit owns a valuable asset, whether the owner is a private capitalist or the government of an underdeveloped country. The asset is worth keeping only if at the margin it earns a return equal to that earned on other kinds of assets. A factory produces things each year of its life, but a mineral deposit just lies there: its owner can realize a return only if he either mines the deposit or if it *increases in value.* So if you are sitting on your little pile of X and confidently expect to be able to sell it for a very high price in the year 2000 because it will be very scarce by then, you must be earning your 5 percent a year, or 10 percent a year, or whatever the going rate of return is, each year between now and 2000. The only way this can happen is for the value of X to go up by 5 percent a year or 10 percent a year. And that means that anyone who wants to use any X any time between now and 2000 will have to pay a price for it that is rising at that same 5 percent or 10 percent a year. Well, it's not happening. Of course, we are exploiting our hoard of exhaustible resources; we have no choice about that. We are certainly exploiting it wastefully, in the sense that we allow each other to dump waste products into the environment without full accounting for costs.

But there is very little evidence that we are exploiting it too fast.

Crowding on planet earth

I have less to say about the question of population growth, because it doesn't seem to involve any difficult conceptual problems. At any time, in any place, there is presumably an optimal size of population—with the property that the average person would be somewhat worse off if the population were a bit larger, and also worse off if the population were a bit smaller. In any real case it must be very difficult to know what the optimum population is, especially because it will change over time as technology changes, and also because it is probably more like a band or zone than a sharply defined number. I mean that if you could somehow plot a graph of economic welfare per person against population size, there would be a very gentle dome or plateau at the top, rather than a sharp peak.

I don't intend to guess what the optimal population for the United States may be. But I am prepared to hazard the guess that there is no point in opting for a perceptibly larger population than we now have, and we might well be content with a slightly smaller one. (I want to emphasize the likelihood that a 15 percent larger or 15 percent smaller population would make very little difference in our standard of well-being. I also want to emphasize that I am talking only about our own country. The underdeveloped world offers very special problems.) My general reason for believing that we should not want a substantially larger population is this. We all know the bad consequences of too large a population: crowding, congestion, excessive pollution, the disappearance of open space—that is why the curve of average well-being eventually turns down at large population sizes. Why does the curve ever climb to a peak in the first place? The generic reason is because of what economists call economies of scale, because it takes a population of a certain size and density to support an efficient chemical industry, or publishing industry, or symphony orchestra, or engineering university, or airline, or computer hardware and software industry, especially if you would like several firms in each, so that they can be partially regulated by their own competition. But after all, it only takes a population of a *certain* size or density to get the benefit of these economies of scale. And I'm prepared to guess that the U.S. economy is already big enough to do so; I find it hard to believe that sheer efficiency would be much served in the United States by having a larger market.

As it happens, recent figures seem to show that the United States is heading for a stationary population: that is to say, the current generation of parents seems to be establishing fertility patterns that will, if continued, cause the population to stabilize some time during the next century. Even so, the absolute size of the population will increase for a while, and level off higher than it is now, because decades of population growth have left us with a bulge of population in the childbearing ages. But I have already argued that a few million more or less hardly make a difference; and a population that has once stabilized might actually decrease, if that came to seem desirable.

At the present moment, at least for the United States, the danger of rapid population growth seems to be the wrong thing to worry about. The main object of public policy in this field ought to be to ensure that the choice of family size is truly a voluntary choice, that access to the best birth-control methods be made universal. That seems to be all that is needed. Of course, we know very little about what governs voluntary fertility, about why the typical notion of a good family size changes from generation to generation. So it is certainly possible that these recent developments will reverse themselves and that population control will again appear on the agenda of public policy. This remains to be seen.

In all this I have said nothing about the Doomsday Models because there is practically nothing that needs to be said. So far as we can tell, they make one very bad mistake: in the face of reason, common sense, and systematic evidence, they seem to assume that at high standards of living, people want more children as they become more affluent (though over most of the observed range, a higher standard of living goes along with smaller families). That error is certainly a serious one in terms of the recent American data—but perhaps it explains why some friends of mine were able to report that they had run a version of the Forrester World Dynamics Model starting with a population of two people and discovered that it blew up in 500 years. Apart from placing the date of the Garden of Eden

in the fifteenth century, what else is new?

There is another analytical error in the models, as Fred Singer has pointed out. Suppose resource exhaustion or increased pollution conspires to bring a reduction in industrial production. The model then says that birth rates will rise because, in the past, low industrial output has been associated with high birth rates. But there is nothing in historical evidence to suggest that a once-rich country will go *back* to high birth rates if (as I doubt will happen) its standard of living falls from an accustomed high level. Common sense suggests that a society in such a position would fight to preserve its standard of living by reducing the desired family size. In any case, this is another example of a poorly founded—or unfounded—assumption introduced to support the likelihood of overshoot-and-collapse.

Paying for pollution

Resource exhaustion and overpopulation: that leaves pollution as the last of the Doomsday Devils. The subject is worth a whole lecture in itself, because it is one of those problems about which economists actually have something important to say to the world, not just to each other. But I must be brief. Fine print aside, I think that what one gets from the Doomsday literature is the notion that air and water and noise pollution are an inescapable accompaniment of economic growth, especially industrial growth. If that is true, then to be against pollution is to be against growth. I realize that in putting the matter so crudely I have been unjust; nevertheless, that is the message that comes across. I think that way of looking at the pollution problem is wrong.

A correct analysis goes something like this. Excessive pollution and degradation of the environment certainly accompany industrial growth and the increasing population density that goes with it. But they are by no means an inescapable by-product. Excessive pollution happens because of an important flaw in the price system. Factories, power plants, municipal sewers, drivers of cars, strip-miners of coal and deep-miners of coal, and all sorts of generators of waste are allowed to dump that waste into the environment, into the atmosphere and into running water and the oceans, without paying the full cost of what they do. No wonder they do too much. So would you, and so

would I. In fact, we actually do—directly as drivers of cars, indirectly as we buy some products at a price which is lower than it ought to be because the producer is not required to pay for using the environment to carry away his wastes, and even more indirectly as we buy things that are made with things that pollute the environment.

This flaw in the price system exists because a scarce resource (the waste-disposal capacity of the environment) goes unpriced; and that happens because it is owned by all of us, as it should be. The flaw can be corrected, either by the simple expedient of regulating the discharge of wastes to the environment by direct control or by the slightly more complicated device of charging special prices—user taxes—to those who dispose of wastes in air or water. These effluent charges do three things: they make pollution-intensive goods expensive, and so reduce the consumption of them; they make pollution-intensive methods of production costly, and so promote abatement of pollution by producers; they generate revenue that can, if desired, be used for the further purification of air or water or for other environmental improvements. Most economists prefer this device of effluent charges to regulation by direct order. This is more than an occupational peculiarity. Use of the price system has certain advantages in efficiency and decentralization. Imposing a physical limit on, say, sulfur dioxide emission is, after all, a little peculiar. It says that you may do so much of a bad thing and pay nothing for the privilege, but after that, the price is infinite. Not surprisingly, one can find a more efficient schedule of pollution abatement through a more sensitive tax schedule.

But this difference of opinion is minor compared with the larger point that needs to be made. The annual cost that would be necessary to meet decent pollution-abatement standards by the end of the century is large, but not staggering. One estimate says that in 1970 we spent about $8.5 billion (in 1967 prices), or 1 percent of GNP, for pollution abatement. An active pollution abatement policy would cost perhaps $50 billion a year by 2000, which would be about 2 percent of GNP by then. That is a small investment of resources: you can see how small it is when you consider that GNP grows by 4 percent or so every year, on the average. Cleaning up air and water would entail a cost that would be a bit like losing one-half of one year's growth, between now and the year 2000. What stands between us and a decent environment is not the curse of industrialization, not an unbearable burden of cost, but just the need to organize ourselves consciously to do some simple and knowable things. Compared with the possibility of an active abatement policy, the policy of stopping economic growth in order to stop pollution would be incredibly inefficient. It would not actually accomplish much, because one really wants to reduce the amount of, say, hydrocarbon emission to a third or a half of *what it is now.* And what no-growth would accomplish, it would do by cutting off your face to spite your nose.

The end of the world—a matter of timing

In the end, that is really my complaint about the Doomsday school. It diverts attention from the really important things that can actually be done, step by step, to make things better. The end of the world *is* at hand—the earth, if you take the long view, will fall into the sun in a few billion years anyway, unless some other disaster happens first. In the meantime, I think we'd be better off passing a strong sulfur-emissions tax, or getting some Highway Trust Fund money allocated to mass transit, or building a humane and decent floor under family incomes, or overriding President Nixon's veto of a strong Water Quality Act, or reforming the tax system, or fending off starvation in Bengal—instead of worrying about the generalized "predicament of mankind."