

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: The Rate and Direction of Inventive Activity: Economic and Social Factors

Volume Author/Editor: Universities-National Bureau Committee for Economic Research, Committee on Economic Growth of the Social Science Research Council

Volume Publisher: Princeton University Press

Volume ISBN: 0-87014-304-2

Volume URL: <http://www.nber.org/books/univ62-1>

Publication Date: 1962

Chapter Title: The Decision Making Problem in Development

Chapter Author: Burton Klein

Chapter URL: <http://www.nber.org/chapters/c2139>

Chapter pages in book: (p. 477 - 508)

The Decision Making Problem in Development

BURTON H. KLEIN

THE RAND CORPORATION

CONTINUING the theme introduced by my colleagues, Marshall and Meckling, I shall attempt in the first part of this paper to describe the nature of the uncertainties involved in development decisions. In the second part I shall turn more explicitly to a discussion of the decision making problem. My purpose is to point out in what key respects the decision making problem in development—especially in areas involving the rather ambitious advances that military development projects typically do involve—is different from that in production. Economists have said a good deal about decision making under uncertainty, but very little about decision making when the responsible person is in a position to do something about the uncertainties facing him. Finally, I shall briefly point out the similarity between the conclusions of this paper and an idea long ago advanced by Schumpeter.

The evidence on which this paper is based is in part the same described in the Marshall-Meckling paper, but it mainly comes from about fifty case studies we have made on military research and development projects, and a few we and others have made on nonmilitary projects (see, for example, Thomas Marschak's paper in this volume, "Strategy and Organization in a System Development Project"). As a basis for any generalizations we would like to have a much larger body of evidence, but two things prevented us from conducting this project as a statistical investigation. For one thing, gathering data on development programs is an extremely expensive kind of research. To find out the cost of development alone can sometimes take months of digging. For another, although we have started some work along these lines, it was not until we were well into the case studies that we began to see that there were some hypotheses that could be statistically tested.

But granting that it would be desirable to have a good deal more evidence, I think some fairly definite conclusions can be drawn from the work we have done to date. Moreover, I believe that the general conclusions of this paper are quite as relevant to civilian R and D as they are to military R and D, providing that in either case we are speaking of kinds of development aimed at ambitious advances.

*Nature of Uncertainties Facing the
Decision Maker*

Before building a bridge, engineers carefully examine the physical environment in which it might be placed, take into account the kind and amount of traffic that will pass over it, the stresses to which it will be subject, and so forth. If the bridge is well designed to begin with, the probability is very high that it will not have to be extensively redesigned during its construction and that it will hold up about as expected. To be sure, unexpected things sometimes happen to bridges after they are built—the initial calculations sometimes prove to be wrong—but I have the impression that such events are rare.

In contrast, development involving the advances that military projects typically seek seldom goes according to plan. No matter how much attention is given to the initial design work, as development goes forward there are almost always a number of surprises. So far as space development projects are concerned, this much can be deduced from reading the newspaper accounts. But I do not think it is generally appreciated how uncertain a business development can be even without getting into space. In almost all the cases we looked into, the system that came out of development was quite different from the one initially conceived (e.g., as different as a Nash and a Buick), and often the system's actual role turned out to be quite different from its intended role. To illustrate this point let me refer to a study we did of six relatively recent fighter-plane development projects. All of these planes were designed for some particular mission, e.g., all-weather interception, ground support. All the aircraft manufacturers based their airframe designs on some particular engine design furnished by one of the engine manufacturers. In almost all these cases, there were also programs for developing specialized electronic and other kinds of equipment. To what extent did these plans materialize? Four out of the six planes ended up with different engines, three with different electronic systems.¹ In order to make them satisfactory flying

¹ Other weapon systems, as well as fighters, tend to come out of development with technological inputs different from those originally planned for them. For example, in looking into the development of 24 weapon-system projects initiated since World War II, we found only three programs in which the final engine coincided with the initial choice. Nine of these programs had three or more engines associated with them at one time or another, several, as many as five. Nor have the British been especially successful in making engine projects come out as planned. Although several British companies develop and produce both airframes and aircraft engines, and although the integrated companies have usually planned their aircraft around their own engines, these companies are seldom able to mate their own engines to their own airframes.

machines, five of the airframes had to be extensively modified; three of the fighters came out of development essentially different airplanes. Of the six planes, three ended up by having quite different operational roles from what was originally planned for them. Only one of the planes possessed the same technological ingredients and had the same kind of operational role that had been initially planned for it. This plane, however, will have a much less important role than was intended for it, in part because another fighter, whose development was started for a different kind of role, has already provided quite as good a capability.

The essential distinction between bridge building and development of fighter airplanes is that whereas the former is mainly a problem of how to take best advantage of existing information, the latter involves also a good deal of learning. In the one case the problem is how to make the best use of known alternatives; in the other, it is to find "good" alternatives. Later in this paper I will discuss the significance this difference has for the decision making problem in research and development. But before turning to this central matter, I want to say something more about the kinds of uncertainties involved in predicting the outcome of development projects.

These uncertainties fall into two broad classes: uncertainties about the cost and availability of various alternatives; and uncertainties about their future usefulness. As Marshall and Meckling have indicated, the utility of military hardware is not highly predictable. Between the time the development of any piece of military apparatus is started and the time it is ready to go into operational use—usually a period of some five to ten years—a host of things can happen to drastically change the demand for it. But it is not only the difficulty of predicting future market situations that makes the development problem unique. Future market situations can be very difficult to predict whether the market be civilian or military, whether the item in question involves a large or a nominal technical advance. Even in the case of bridge building, it may be assumed that predicting future demand is no trifling problem.

On the other hand, what does sharply differentiate activities like bridge building from activities like fighter plane development is that the latter is also characterized by considerable uncertainties on the supply side. In the one case the decision maker can proceed with a considerable degree of confidence that the information on technical characteristics, costs, and time to complete the project available to

him at the beginning of the project will not significantly change; in the other he has to proceed under the assumption that it probably will change.

Why is it so difficult to foretell how alternative designs will turn out? In some instances discoveries have been made in development like the discoveries that come out of basic research. But even with no scientific discoveries made during the course of development, predictions often have turned out to be bad. As an illustration, consider the postwar experience of this country in airborne radar development. Almost all of the radars that have been developed since World War II are based on principles that were quite well known at the end of the war; in fact, the basic designs of many of the radars developed in the last fifteen years were conceived during the war. Yet it is very difficult to find any radar that, when a seemingly completed system was first tested, met its original expectation. Out of the dozen programs we examined, there were only two in which the development costs after a full-scale test of the radar did not exceed by several times the costs up to the time of such a test. And there were none in which the cost of meeting the unanticipated problems was negligible. Thus, even though the problems were not of a scientific nature, it can hardly be argued that a good deal was not learned during these programs.

What was learned? The learning that took place in these radar programs, or for that matter the learning that takes place in almost any kind of development program, can be roughly divided into three stages. In the earlier stages of a program the dominant problems are usually those associated with overcoming the major technological problems. The seriousness of these problems depends in good part on how ambitious the advances are that are being sought. In military research and development, even when the advance sought is relatively modest, a variety of technical problems almost invariably turns up. In the second stage of development the main problems have to do with getting the various pieces of a system to perform with a tolerable degree of reliability. In radar development, probably much more has been spent on overcoming reliability problems than on all other problems combined. As might be deduced from reading the newspaper accounts, it also has been a problem in ballistic missile development. Finally, there are the problems involved in mating the various subsystems, and in tailoring the equipment to work satisfactorily in the environment it will be used in, and in finding efficient means for getting it produced.

As one might suspect, in actual practice the learning process does not always follow this precise chronology. For example, in the initial design stage the developer will be concerned with problems of reliability and of how to combine various technical inputs. He knows from experience that some types of designs pose an inherently more difficult reliability problem than others; he also knows that some kinds of constraints are required to insure that the various parts of the system will be ultimately compatible. Conversely, he may find himself confronted with serious technical problems very late in the program. This may be because the development program was not conducted in a manner to uncover them early; or it may be because some of the "later" problems, e.g., reliability, turned out to be so serious that it was necessary to redesign the equipment, which, in turn, gave rise to a new series of technical problems.

Without implying that the three stages of learning, as I have outlined them, are in practice easily distinguished from each other, let me now try to suggest the nature of the problems that do come up.

Designing a compressor for a jet engine provides a fairly typical example of the kind of technological problem the developer has to solve. In a number of important respects the performance of an engine is dependent on the design of the compressor; hence the continuing emphasis on getting improved compressor designs. The root of the problem in compressor development is that, although on the basis of the fundamental equations of aerodynamics it is "theoretically" possible to predict the performance of any particular design, in actual fact it is possible to predict only in very approximate terms the characteristics required to provide a given performance. And in fact while some of the particular designs that satisfy the criteria of a "good" design may turn out to be good; others may turn out to be bad. This being the case, the chance of any design turning out to be a good one will depend partly on the experience and ingenuity of the engineers and partly on luck. But in any case, when a good design is finally hit upon, the outcome usually has been preceded by a tremendous amount of experimental work.

It can be argued, of course, that the basic reason for this kind of situation lies with science—that if the scientific basis were more complete, better and more precise predictions could be made. This I do not dispute. My point is simply that science cannot usually predict how a particular design will turn out. It is true that there are significant differences among technologies as to the kinds of predictions that can

be made. In some branches of nuclear technology it has been possible to make quite accurate predictions. On the other hand, in other technologies—notably materials—the scientific basis for making any kind of prediction is very meager, with the result that progress is almost entirely made by trial and error.

Experience indicates that one of the most unpredictable things about a new piece of equipment is the problems that will be encountered in making it perform reliably, and the time that it will take to overcome these problems. One of the main reasons why development costs tend to be seriously underestimated is a tendency to minimize the reliability problems that will be encountered and the amount of testing and modification that will be required to overcome them. Quite apart from all the other problems, the simple problem of diagnosis has often proved to be a formidable one. Experts evaluating a test program have often not been able to agree among themselves as to the cause of a major malfunction. The process by which reliability problems are overcome can hardly be described as anything but a highly empirical one: a piece of equipment is run until it breaks down (which in the early stages of the program is not very long); a series of modifications is undertaken to correct the major causes of malfunction; and when the equipment can perform reliably under relatively easy conditions, it is then subject to more and more stringent conditions. In some cases the result of this process is a host of minor modifications; in others, it is some fairly major changes in the design of the equipment.

Before reliability work on the various parts of a major system has progressed very far, another major problem will have come to the forefront: how the various elements of the system can best be combined to perform some useful function. As indicated earlier, this problem certainly will have been considered in the initial design work, but in order that it can be considered very fruitfully, and in detail, development must have progressed to a point where a good deal is known about the kinds of trade-offs that are actually possible. These may differ considerably from the kinds of trade-offs initially believed possible. For example, whereas the weight of a guidance system of some specified characteristics may have been initially estimated at, say, 300 pounds, its actual weight may turn out to be four times that amount. And this in turn may force some major changes in other parts of the system. It may be found that in order to get a fire-control system working with an acceptable degree of reliability, some of its

versatility has to be given up and the loss in capability partly made up in some other way. Or, as often happens, it may simply be found that a major component will not be available until much later than the rest of the system, and that some alternative device of quite different characteristics has to be substituted.

As development progresses a good deal is learned not only about the technological alternatives themselves, but also about the environment in which the system will operate. Almost invariably, the actual conditions under which a new piece of equipment will have to operate are found to be quite different from the assumed conditions. For example, after only a little experience in testing a missile it may be found that the kind of ground equipment needed to support it is of a quite different kind than was earlier assumed. It may be found that the design of the equipment requires far more skilled personnel to operate and maintain it than are available. Time after time, equipment has to be modified so that a large corps of engineers and scientists will not be required to maintain it.

Finally, it often turns out that new methods have to be developed in getting the equipment into production. Sometimes the problems are so basic as to require a good deal of research and development work on the production processes themselves. For example, in connection with developing the B-58 supersonic bomber a process for producing honeycombed aluminum structures was developed. Work on the development of the B-70 supersonic bomber will involve the pioneering of techniques for producing stainless-steel structures. The increasing emphasis that has been placed on the development of miniature devices of one kind or another has resulted in a number of new production problems.

Though sometimes there are serious problems involved in getting new kinds of equipment or devices produced, the evidence we have examined suggests those problems are usually much easier to overcome than the problems involved in their development. Ordinarily, comparatively little effort is required to overcome the production problems.

Though I might have gone further in characterizing the kinds of problems that come up in development, I trust enough has been said to indicate that in the course of development a number of different kinds of problems have to be solved—that development is essentially a learning process. Its main purpose is to find out how much new kinds of capabilities will cost, and how useful they can be made. As Meckling

and Marshall point out in their paper, the very large uncertainties associated with cost estimates of not-yet-developed systems are mainly uncertainties about what it is that is being costed. These are the uncertainties I have been discussing.

But granted that a good deal of learning is involved in development, to what extent is the problem of making choices here essentially different from what it is in the familiar cases in which perfect knowledge is assumed from the outset? Granted that there is often a good deal more risk in making development decisions than other kinds, is the decision making problem here any different, essentially, from what it is in other less risky areas? To this question I now turn.

Central Differences between the Decision Making Problem in Development and in Production


DECISION MAKING STRATEGY

The fact that it is very difficult to make good estimates either of cost or demand does not in itself prove that the decision making problem in development is basically different from other kinds of decision making problems that involve choices among various kinds of alternatives. One might simply take the attitude that it makes no difference what kind of probability distribution is associated with the initial estimates; that the decision maker has no alternative but to make the best estimates he can, and to act as if he were dealing with certainties.² This situation is roughly illustrated by case A on Chart 1. The decision maker, we assume, is trying to decide whether to develop device A or device B, and that his decision will be entirely based on what he believes these devices will ultimately cost. It can be seen that in case A the initial subjective probability distributions are very wide, and that they remain very wide to the end of the project. Providing he has to make a choice, it is apparent that in this situation the decision maker will choose the device whose expected cost (the mean of the subjective distribution) is lower—in this example, device B. On the other hand, assume that the distributions can be significantly narrowed long before the end of the project—that the situation is more like that shown

² To be sure, under such circumstances the decision maker (if he could afford it) might decide to insure himself against future uncertainties by developing two airplanes instead of one. But the idea of taking contingencies into account in planning, while certainly an important aspect of the problem of making decisions under conditions of uncertainty, is a notion long familiar to economists.

CHART 1


Variance of Estimates as a Function of the Stage of Development



in case B. Then it is by no means apparent that, even though he must ultimately choose either device A or B, his best strategy will be to choose that device with the lower initial expected cost. The crucial question is the extent to which he is in a position to reduce these uncertainties, and at what cost.

To push this argument further, let us assume that the developer starts out with estimates of performance, cost, etc., whose accuracy is represented by point A on Chart 2. When development is completed,

CHART 2
Rate of Learning Hypotheses



the accuracy of the estimates will have improved to B. If the actual world corresponds more or less to Y—if there are large indivisibilities in this learning process—then, although the decision maker may feel very uncomfortable about his initial estimates, he has only one thing to worry about: should he become involved in such a risky undertaking or should he not? If he does decide to go ahead with the project, he may as well act on the assumption that his initial estimates were good, for there is no way to find out except by committing himself to the

entire project. If it costs \$100 million to complete the development of the system in question, there is no way he can significantly improve his initial estimates except by spending the entire \$100 million. It does not matter whether in choosing this particular alternative he happens to regard the initial probability distributions of cost, availability, and demand as highly compact or highly dispersed; he has to make choices on the basis of expected values. In other words, he has to proceed as he would in dealing with a conventional kind of optimization problem.

On the other hand, if the actual possibilities for improving the estimates lie in the neighborhood of the line AC, the decision maker may want to pursue a very different kind of strategy. Under these circumstances he can significantly narrow the uncertainties facing him long before the end of the program. His decision whether or not to adopt such a course of action will depend essentially on what it is likely to cost. If it is relatively inexpensive (as compared with total development costs) to carry projects into the initial phases of development—if the initial uncertainties can be significantly narrowed for a cost of not \$100 million but, say, \$15 or \$20 million—then it is fairly apparent that the decision maker will want to play the game very differently. Under these circumstances, the basic nature of the decision problem will be much more like that involved in solving a crime than that involved in solving the typical kind of optimization problem. The essential problem for the developer is how to eliminate potential suspects quickly and efficiently.

Since it is a matter of such key importance, we have been interested in finding out all we can about the opportunities, as they actually exist, for improving information as development progresses. The general impression we have obtained from our case studies is as follows. In the first place, as might be expected, we have found that there are considerable differences among development programs. For example, it is quite obvious that if the advances sought are relatively modest (e.g., as in developing a transport plane), the uncertainties are ordinarily much more easily resolved than if the advances sought are relatively ambitious (e.g., as in developing a new kind of missile). The possibilities for early learning will also depend on the kind of technology involved. Pursuing multiple approaches to difficult technological problems is a good deal cheaper in some fields than in others: in relation to total development costs, the cost of getting preliminary test results is much cheaper in developing missile guidance

systems than it is in developing large solid-propellant missile engines. Finally, the rate of learning will depend also on the spirit in which the development program is conducted—when the developer is very keen to learn, he usually finds ways to make the learning process more divisible.

But, though there are considerable differences among development programs, we feel that cases involving very large indivisibilities in learning—cases approaching Y—are in fact rare. And while it is often true that some important problems are not resolved until late in the development process (and sometimes never), it is also often true that very modest expenditures can buy a considerable amount of information. Consider, for example, the experience in ballistic missile development. As might be expected, it has been found that there is no way of predicting exactly how an entire system will work short of assembling it and testing it a number of times. On the other hand, it also has been found that a surprisingly large amount of information can be obtained by relatively inexpensive kinds of development work and testing that can be done on the ground, for example, static test firings of engines or sled tests of guidance systems. Even after the test firings have begun, it usually is not efficient to test an entire system until a whole sequence of preliminary tests has been conducted. Or take the problem of developing a jet engine. Here again serious unanticipated problems may crop up late in development—after more than, say, \$150 million has been spent. But experience has also indicated that the test of a preliminary model aimed at a demonstration of its novel components, costing perhaps \$5 to \$15 million, can make possible much better predictions of an engine's ultimate performance than predictions based simply on a design study. And by the time roughly one-third to one-half of the entire development has been paid for, it usually has become fairly apparent which engines are going to turn out to be good and which are not.

Fairly recently we began some work aimed at more rigorously testing the "early learning" hypothesis. The general idea behind this study is to find out how rapidly estimates of availability, cost, and performance improve as a function of development time and development cost. If the accuracy of the estimates is found to improve at a linear or increasing rate, we would regard this as confirmation of the hypothesis. To date, our efforts along these lines have consisted of a fairly crude analysis of a small number of availability estimates, and a much more sophisticated analysis of a somewhat larger body of

cost estimates. Though the second study has not been completed, I was able to persuade Robert Summers to throw caution to the winds, and provide me with some of his preliminary results.

Our analysis of availability estimates—which was undertaken more or less as a pilot investigation—was restricted to 8 missile programs, for which there were 35 estimates of availability made at various stages of the programs. In order to get an absolute measure of the error in the estimates, the estimated time to go (ETG) was subtracted from the actual time to go (ATG). To obtain a relative measure, the error rate was computed as $\frac{ATG - ETG}{ATG}$. On the basis of this concept, an


error of five years in an estimate made ten years from the actual time of completion would be equivalent to an error of one year in an estimate made two years from the time of completion (0.5 in either case).

As might be expected, there was a fairly substantial difference among the programs with regard to the rate of improvement in the estimates. In one of the programs, the predictions made after development was started actually turned out to be a good deal worse than the initial prediction, and there was no significant improvement in the estimates until rather late in the program. However, in all the other cases, improvement in the estimates was fairly continuous; in five of them the error rate declined very rapidly. The adjusted errors³ as a function of development time are plotted on Chart 3. As the chart indicates, the errors fell quite steadily as development progressed. Plotting the error rate against development time would indicate an approximately constant rate of improvement in the estimates.

Roughly the same conclusion—improvement in the accuracy of the estimates at about a constant rate—is indicated by the preliminary results of our analysis of cost estimates. These results are based on 22 weapon systems, both aircraft and missile, for which we had 71 usable observations. As in the availability study, the weapons chosen were those on which we were able to get information (though at no small cost) and, though we obviously cannot be sure, we have no reason to believe that the errors made in costing these systems were any larger than the errors that have been made in costing other aircraft and missile systems.

³ Since the number of observations was relatively few and since we are here primarily interested in comparing the manner the errors decline rather than comparing the differences in the absolute errors among various programs, all of the errors were deflated by dividing them by the interpolated values for the mid-points of the programs.

CHART 3
Errors in Availability Estimates



In order to test the early learning hypothesis the data were classified in two ways: according to the stage of development at which the estimate was made, and according to the ambitiousness of the advance sought in the program.⁴ On the basis of the hypothesis, we should expect the variance in the cost estimates to decline in relation to the stage of development at which the estimate was made and to increase in relation to the degree of advance being sought.

The results of the investigation are summarized in Chart 4. Shown in each cell of the table are the summary numbers for the average errors $\frac{\text{(actual cost)}}{\text{(estimated cost)}}$ and the standard deviations and the number of observations on which they were based. The arrows indicate whether the direction of the changes observed in the average errors

⁴ In order to decide whether the jump in the state-of-the-arts was "small," "medium," or "large," a panel of engineers was individually polled. Though in the case of a few programs there were considerable differences of opinion, on the whole the judgments on the degree of advance involved were fairly uniform.

DECISION MAKING PROBLEM IN DEVELOPMENT

CHART 4
Technological Advance

$$F = \frac{\text{Actual cost}}{\text{Estimated cost}}$$

	Small	Medium	Big
Time	Early $\bar{F} = 0.80$ $\hat{\sigma} = 0$ (1)	$\bar{F} = 2.15$ $\hat{\sigma} = 0.57$ (4)	$\bar{F} = 3.54$ $\hat{\sigma} = 2.17$ (3)
	Middle $\bar{F} = 1.50$ $\hat{\sigma} = 0.6$ (2)	$\bar{F} = 1.32$ $\hat{\sigma} = 0.39$ (6)	$\bar{F} = 2.52$ $\hat{\sigma} = 1.41$ (4)
	Late $\bar{F} = 1.11$ $\hat{\sigma} = 0.34$ (8)	$\bar{F} = 1.06$ $\hat{\sigma} = 0.18$ (8)	$\bar{F} = 1.16$ $\hat{\sigma} = 0.34$ (5)

and standard deviations, in looking across the rows and down the columns, is consonant or at variance with the double hypothesis. The solid arrows run in the direction predicted by hypothesis; the dashed arrows do not.⁵

It can be seen that 9 out of the 12 arrows point in the "right" direction and that 3 do not. However, only the encircled arrows indicate changes that are, roughly speaking, statistically significant—

⁵ Though on the basis of the hypothesis the arrows should generally point to the right and downward, actually we should not expect to find significant differences among the bottom three rows. The accuracy of cost estimates made late in the development process should not be very different whether the particular project happened to involve a big advance or a small advance.

and it can be noted that all of these point either to the right or downward. As might be expected, although they are pretty much the same for the late time period, for the earlier two periods the size of the standard deviations appears to be highly correlated with the size of the technological advance being sought. And looking down the columns, the change in the standard deviation as well as the average error appears to be roughly linear. By taking a few liberties with the data, it might even be argued that the evidence is consistent with the assumption of a declining rate of error, but as someone once observed, "There is a limit to the amount of blood that can be squeezed out of a turnip."

It will be noted that I have expressed the errors in the availability and cost estimates as a function of development time rather than of development expenditures. Though it would have been more pertinent to make comparisons in terms of development expenditures, unfortunately we do not have that data for all the programs included. However, had it been possible to relate improvements in the accuracy of the estimates to development expenditures, there is little doubt that the results would have been somewhat more favorable to the hypothesis advanced earlier; that is, there is little doubt that as a function of development expenditures the accuracy of the estimates would have improved more rapidly. This is because the early stages of development are relatively inexpensive as compared with the later stages. We have found that by the time one-half of the development period has elapsed, the proportion of the total development bill that will have been spent is typically a good deal less than one-half.

To be sure, it would be desirable to have much more evidence on the rate that knowledge increases as development progresses. And we do hope that we will be able to push this kind of inquiry further, mainly because we believe it might provide some deeper insights into the nature of development. But I doubt that more evidence would sharply upset the proposition that, as development progresses, there is likely to be a more or less steady improvement in knowledge.

If the information used for making subsequent decisions can be markedly improved as development proceeds, progress in development should be much more rapid when the decision maker deliberately chooses a sequential decision making strategy. In comparing the results of the various development programs we studied, we think we have observed that there is a close relationship between the degree of success achieved and the willingness of the developer to

undertake the project with the conviction that before it is finished he will have much to learn. In almost all the outstanding programs we have examined, the developer has started out with a very loose definition of the system he was trying to develop and has exhibited a considerable willingness to get hardware—even in rudimentary form—into test early in the program, and to pursue multiple approaches to difficult technical problems.

As a result of the work we have done, we are also coming to suspect that those firms that have been successful in research and development over a long period employ a kind of development strategy very different from that of firms whose success has been much more limited. For example, we have found that the “personalities” of the two firms that have been outstanding in jet engine development—Pratt and Whitney in this country, and Rolls Royce in Britain—are much alike, and that their approach in development tends to be a good deal more pragmatic than that of their competitors: “You develop an engine by testing it, and you never develop an engine with only one airplane in mind.” In order to find out further whether there really is anything in the idea that the more successful firms in R and D go about the business differently from the less successful firms, we have initiated a project aimed at finding out as much as we can about the kinds of approaches private firms employ in research and development. We began with the highly successful firms (for the good reason that they were inclined to be much more cooperative in furnishing us information), and the results of our first investigation are reported in Thomas Marschak’s paper “Strategy and Organization in a Systems Development Project.” As it turned out, the manner in which Bell Telephone Laboratories went about developing a microwave relay system provides a first-rate example of sequential decision making.

EFFICIENCY IN DEVELOPMENT


If it is granted that the kind of strategy required to play the game of development is very different from that required to play the game of production (or other kinds of games in which the initial uncertainties have to be taken as given), the next question I would like to raise is this: are they really the same kind of games? I assert that they are not, and that the kind of behavior that will be efficient in one will not be in the other.

To make it clearer what I mean, let me take the problem of choosing an engine for an aircraft and ask, first, how the fellow who looks at

the problem as a production decision would go about making the choice. Obviously, the first question he wants an answer to is, what kind of airplane?—for the engine characteristics well suited to one kind of plane may not be well suited to another. In bombers and transport planes, for example, fuel economy is a much more important factor relative to engine weight than it is in fighters, because in the former the weight of the fuel to be carried is very large in comparison with the weight of the engine. Moreover, for engines based on current technology these characteristics are highly substitutable.

Given this information, it is fairly obvious how the producer will make his decision. If I on Chart 5 represents the trade-off curve based on current technology, and if the plane in question is a bomber, he

CHART 5
Aircraft Engine Alternatives



I — Current technology
II, III, IV — Advanced technology

will order engine B, and delivery of an engine possessing that particular combination of characteristics will meet his objective, on the other hand, if engine F is delivered to him he will be unhappy, for he will not have an optimal airframe-engine combination.

So far as the developer is concerned, he is primarily interested in developing a much better engine than any represented on I. His principal concern is whether he is likely to end up on a curve only slightly higher than can be provided by present technology (e.g., II) or a much higher one (e.g., IV) that will dominate an optimized present engine in all important respects. But within fairly wide limits, which particular point he finally manages to reach is of rather secondary importance to him.

Let us suppose, however, that the developer does make up his mind that he is going to have his cake and eat it, and that he sets for himself the firm objective of getting to point B on III. That he will be able to get there eventually there is little question. Given enough time and effort in development almost anything is possible. What I do question is whether in doing so he will have acted efficiently. With the same effort, I contend, the developer almost certainly could have gotten to some point on a higher opportunities curve, e.g., on IV. My proposition, in other words, is that efficiency in the usual sense of the word can be attained only after considerable sacrifice in progress, i.e., reaching a higher opportunities curve.

Unfortunately, in this paper I cannot provide the sort of evidence necessary to demonstrate to the skeptic that this is a reasonable proposition. But I would like to indicate briefly why I believe that it is.

One of the reasons—already pointed out—is that development projects tend not to come out as they were initially planned. For example, when we ask what the experience has been in aircraft development programs we find that airframes seldom do end up with the engines planned for them. As a matter of fact, engines that win the competitions, no matter which particular plane they were initially designed for, almost invariably dominate the losers in all important respects and are almost invariably used in a wide variety of planes, from commercial airplanes to supersonic fighters. In the case of missiles, it is astonishing how often successful systems are developed on the basis of technological components not originally intended for them. *Ex post*, a missile system can be reasonably defined as a system mainly made from components developed for other missile systems.

If it is accepted as a fact of life that nothing much can be done to

combat the tendency of development projects not to come out as originally planned, the whole idea of attempting to reach some predetermined optimal point makes little sense. No matter whether an engine with a somewhat different combination of characteristics could have been developed, the best engine at the time the choice actually has to be made is the best one available.

The second reason back of my general proposition is that if it is not accepted as a fact of life that technology is ill mannered, it costs a good deal to force it to be well behaved. While the kind of efficiency practiced in production may be all right in its place, in development it is terribly expensive. When Marshall pointed out that the pain involved in carefully comparing costs with marginal utility can sometimes exceed the benefits, he was probably thinking of some of the calculations an individual might make before purchasing some item of personal consumption. But in development, optimization involves a good deal more than making some laborious calculations. For one thing, it involves putting tight constraints on the engineers who work on the various parts of the system, and, especially if the advances aimed for are ambitious, it enormously complicates their development task. For another, even if a great effort is made to have the individual pieces come out as originally planned, usually they do not; and this seriously complicates the job of coordinating the entire project, which even in the best circumstances is formidable. Hence my belief that the goal in development is not the same kind of efficiency one strives for in production.

Some Closing Observations

In observing that the decision making problem in development and in production are two very different sorts of things, I am merely calling attention to something that Schumpeter pointed out long ago, namely, that the conditions required for an efficient allocation of existing resources are not the same as the conditions required for rapid economic progress.⁶ Out of this argument came his argument for monopolies, for which, certainly in recent years, he has become much more famous. But without wishing to argue the point here, I contend that no matter what one thinks of Schumpeter's views on monopoly, they do not necessarily follow from his central idea, and actually have had the effect of obscuring it. Even though everything

⁶ J. A. Schumpeter, *The Theory of Economic Development*, Harvard University Press, 1911, Chap. II.

else he said may have been wrong, for this idea alone Schumpeter deserves a high place on the list of distinguished economists.

In 1909, when Schumpeter was writing his *Theory of Economic Development*, he was obviously not thinking of research and development as an organized activity. But the instinct to try out new ideas, as opposed to the instinct to satisfy existing wants efficiently, is of course not confined to research and development laboratories. Nor, obviously, is the experimental approach an approach used only by scientists. In other words, the classical theory of the firm provides a poor description of how the firm actually does behave—that is, the firm that gets ahead rapidly.

On the other hand, it may also be true that the forces making for efficiency in the narrower sense are deeply engrained in our society and stand in the way of more rapid progress. While economists have probably had little influence on business practices in research and development, the same cannot be said of cost accountants, management experts, and the growing army of business school graduates in general. And to their influence must be added the influence of the engineers. Despite the rise of experimental physics, a good many engineers are still trained in the tradition that it is sinful to design anything that might later have to be changed.

COMMENT

FREDERIC M. SCHERER, Harvard University Weapons Acquisition Research Project

My comments are addressed to Burton Klein's paper, "The Decision Making Problem in Development," and to the Marshall-Meckling paper which provides data upon which it is based. Specifically, I should like to make several observations on the nature of uncertainty in development decisions and on the possibilities for doing something about such uncertainty.

The efforts of Klein and his colleagues to determine empirically the rate of learning in development projects are highly commendable. If it can be demonstrated that the uncertainties in a development program can be substantially reduced early in the game (thus, the rate of learning function X in Klein's Chart 2), then the argument for multiple approaches and programs is strongly supported. Nevertheless, I think it necessary to emphasize that there are many different

parameters whose values will be learned during a development effort, and each of these undoubtedly has a different rate of learning function. My own interpretation of data collected by the Harvard Weapons Acquisition Research Project suggests that uncertainties regarding technical feasibility and specific technical performance parameters do in fact decline sharply early in the development. However, accuracy in predicting development time and cost and production cost probably increases at a roughly linear rate over time, while uncertainties about reliability, operating cost, and military utility are usually not reduced substantially until late in the development effort (thus, the rate of learning function Y in Klein's Chart 2). If this is so, then the development approach chosen will depend upon which parameters happen to be particularly critical in the decision. For radical new innovations, learning about technical feasibility and performance may be most important, while for minor improvements, reliability and operating cost considerations may be paramount.

Secondly, I would disagree with any interpretation that the variances shown in the Marshall-Meckling paper are wholly the inevitable result of pressing an uncertain state of the art; that, as Klein suggests, "nothing much can be done to combat the tendency of development projects not to come out as originally planned." Granted, there will be unexpected technical difficulties or pleasant surprises, or both, during almost any advanced development effort. Granted, too, that it would be the height of foolishness not to include provision for unexpected contingencies in development plans, although the overt use of contingencies is usually prohibited by government contracting regulations. Nevertheless, the occurrence of truly unforeseeable and unavoidable technical difficulties represents only one of several variance causes. (By unforeseeable, I mean difficulties that could not have been anticipated by a good hard look at the state of the art—which has not been taken at the outset of all too many military programs.) Analysis of the Harvard group data shows that most of the variances from original time and cost predictions were man-made rather than caused by an ill-natured technology.

One common cause of prediction variances was unrealism in the original estimates, usually because of attempts to "sell" a program. Knowledgeable military decision makers usually apply heavy discount factors to such estimates.

Another major cause of variances was the intentional decision of military officials not to adhere to the original plan. Any develop-

ment effort requires many major and minor compromises between essentially conflicting sets of objective functions—maximization of technical performance values and reliability on the one hand, and minimization of development time, and development, production, and operating costs on the other hand. The choice made in each individual case depends upon the relationship between the military utility of performance increments or time savings and the costs associated with those increments.¹

Sometimes the need for a weapon may be so urgent that the system will be ordered into production as quickly as possible at the sacrifice of some technical perfection. In such cases, development cost and especially time predictions are apt to prove quite accurate, simply because maintaining the schedule was of vital importance. And typically, what is sacrificed is just a few percentage points in relevant performance values such as speed, range, CEP, operating ease, etc. It is commonly said that there is a diminishing returns phenomenon in development; that getting the last 10 per cent of performance or reliability takes 50 per cent of the effort. I have seen some evidence to support that hypothesis, although Klein and his associates might make a major contribution to our understanding of research and development economics by bringing their empirical data to bear on a rigorous test. In any event, attaining an additional ten per cent or so in performance or reliability, or both, may in certain cases be deemed so much more important than early tactical availability of the weapon that development time (and hence cost) will be *allowed* to exceed original predictions.

Let me use Klein's fighter-engine mating example to show what I mean about the man-made nature of these problems, taking one extreme example which is probably included in his sample of six fighters and which has been studied in detail by the Harvard group. This particular fighter was associated with four different engines during its development. The first-choice engine was dropped because it was heavily committed to a higher priority aircraft. The second choice (programmed for interim use only) was dropped when it fell a year behind schedule and because the engine manufacturer was unwilling to produce the very few engines that interim use implied. The third choice was used on a strictly interim basis until a better engine

¹ The theoretical relationship of these factors is discussed in Chapters IX and XVII of M. J. Peck and F. M. Scherer, *The Weapons Acquisition Process: An Economic Analysis*, (to be published in 1962).

became available. The fourth engine was programmed and actually used in the tactical version. Significantly, it was recognized right from the beginning that the fighter might use any of the engines; in other words, the engine question was not so uncertain that the problem could not be anticipated. It seems to me that this kind of situation has been more typical than the case in which an engine (or other subsystem) unexpectedly turns out to be unsuitable, as in the ill-fated F3H-1-J40 marriage.²

Another qualification needs to be made about Klein's six fighter illustration. The greater portion of all six developments took place during a period of relatively little international tension. The Korean conflict (which precipitated numerous performance sacrifices to get fighter aircraft into production quickly) had subsided, and we were not yet worried about the threat of high performance Russian bombers. We were well supplied with medium performance, yet adequate, interceptors and fighter-bombers. Under these circumstances, we could afford to change development plans in midstream in order to get the best possible engines, aerodynamic configurations, and electronic subsystems into our aircraft. Had the Korean conflict erupted into something bigger, we would have looked at such "nice to have" changes in a quite different light.

Other man-made causes of time and cost variances in weapons program outcomes included the lack of funds to implement production goals, bureaucratic decision making delays, and the failure of contractors to bring adequate human resources to bear (either because their organization lacked highly capable scientists and engineers or because those people were assigned to projects with greater promise of long term profits). This does not exhaust the list of causes I have observed, but at least it suggests that all variances from original predictions are not the result of a malevolent nature. If pressed for a quantitative estimate, I would guess that less than 30 per cent of the development time and cost variances experienced in development programs studied by the Harvard group were due to technical difficulties which were both unforeseeable and unavoidable.

If in fact development project outcomes are not nearly so far beyond human control as comparisons of original predictions with actual outcomes would suggest, then rather difficult conclusions about development approach and strategy are indicated. For one, I fail to

² Cf. *Navy Jet Engine Procurement Program*, H.R. Committee on Government Operations, 10th Intermediate Report, 1956.

share Klein's aversion to "putting tight constraints on the engineers who work on the various parts of the system," although our difference in this matter may hinge on semantic difficulties.

In the first place, clear and firm technical constraints are necessary to ensure that interrelated components fit together properly. Imagine what consternation would result from attempting to cram a 12 cubic foot bombing and navigation system into 9 cubic feet of airframe volume, or from feeding digital inputs into an analog computer, or when each of several component developers uses upon his individual item the whole error factor for a guidance system! Moreover, there are good psychological reasons for establishing a technical plan, even if it must later be modified. Most engineers, except perhaps the few really creative ones who ought to be engaged in basic and applied research, need to have a set of specific performance and time goals toward which they can work.

I think what Klein really means is that flexibility is required in developing basically new technical concepts. If so, I agree fully. Clearly, the development of a jet engine incorporating a radically new approach to compressor design should not be tightly constrained, at least not until the design's soundness has been demonstrated. But there comes a time in jet engine development when one must begin worrying about whether or not an afterburner will be used, about the kinds of air intake geometry to be employed, mounting arrangements, and the host of other details that must be solved (usually at a much greater expense in time and money than proving the basic concept required) before the engine can be mated with a specific airplane. When that time arrives, the world of loosely defined breakthroughs is forsaken for the more humdrum schedule, use, and cost oriented world of system development.

My observations about the controllability of development program outcomes also have implications regarding the use of multiple approaches to development problems. If systems development is not really so uncertain, then pursuing multiple approaches may not be as important on the average as Klein suggests. Furthermore, it should be recognized that there are costs other than dollar outlays to be weighed against the risk reduction benefits. One obvious cost is the dissipation of very scarce topflight technical talent. I think there are many cases in which it is preferable to focus all or most of the best talent available on one selected approach rather than to spread it over many efforts.

Another cost is related to the impact of multiple approaches upon

decision making. There are two different processes going on simultaneously in a development effort: the creative process and the decision making or selective process. The creative process generates knowledge in the form of specific technical alternatives by defining problems, proposing solutions, and testing proposals. So long as this process indicates clearly which one alternative solution to any given problem is best, the decision making part is easy. But more frequently, alternatives are generated at a faster rate than best solutions can be identified. Usually the choice among alternatives is not obvious; one offers certain advantages and another, others, and so decision makers must indulge in time consuming evaluations before making their choices. If too liberal a use of multiple approaches is allowed, the decision making process can become overloaded and break down with consequent costly delays. Possibly this can be averted by superior organization, for example, through effective decentralization of decisions. Nevertheless, I have seen instances in which breakdowns did occur when the menu of technical alternatives became too rich for a project management group. Certainly, the theory of the firm has not yet abandoned the idea that the entrepreneurial (decision making) function imposes a limit on the scale of operations.

This does not mean that multiple approaches to development problems should be avoided. It merely means that they should be reserved for certain limited situations, the exact extent of their use depending upon the interrelationship of many factors: the kinds of uncertainties attached to specific problems, the rate of learning at which these uncertainties will be dispelled, the availability of technical talent to man individual projects, the monetary cost of additional approaches, the capacity of the decision making organization, the importance of minimizing development time, the military utility of marginal performance or reliability increments, etc.

One additional complication to the choice of a development strategy should be mentioned. A crucial factor in the decision whether or not to employ multiple technical approaches to a development problem is the probability distribution describing one group's chances of reaching an acceptable (in terms of time and technical parameters) solution. However, the introduction of interfirm competition will usually alter the original single-group probability distribution. Competition may improve the prospects for an acceptable solution if it incites the contractor to assign his best technical talent to the effort. It may impair those prospects if the contractor perceives his

chances of coming up with the winning solution to be so poor relative to alternative opportunities for employing his resources that he allocates his talent into more promising areas.

What I have tried to show in these comments is that defining a single best approach to the development of advanced systems is exceedingly difficult, particularly since the historical evidence contains often obscure cause and effect relationships. This does not mean it should not be attempted, and the RAND group working on the problem should be complimented for its progress in the area. Nevertheless, I am afraid any theory of an optimal development strategy must necessarily be quite complex in order to reflect the vast range of situational factors involved.

REPLY BY KLEIN

Before taking up Frederic Scherer's specific charges, I would like to make this introductory comment: Contrary to the impression he apparently got from reading my piece, I certainly do not believe that the only kinds of uncertainties involved in making development decisions are technical uncertainties. The decision maker is uncertain not only about the kinds of performance that can be achieved, when, and at what cost—but also about what the military utility of particular capabilities will be some five to ten years in the future. Commonly, weapon systems go through a progression of changes during their development because earlier notions as to what can be achieved and what might be most useful *both* change. To be sure, in writing my paper I did give a good deal of attention to the technical uncertainties, but that was simply because it is these uncertainties that give development its special flavor.

Scherer's central criticism of my paper, I take it, is that I have considerably exaggerated the uncertainties in development. As he puts it, most of the variances from the initial predictions—about 70 per cent, he judges—were man-made rather than caused by an ill-natured technology. It would be nice if development were a much more predictable business than it seems to be, but before accepting Scherer's proposition I want a good deal more evidence. As it is, I am not even sure what he means when he says that "most of the variances were man-made." Does he mean that, after considering the nature of what was learned during the course of the development projects the Harvard Group reviewed, he personally has come to the

conclusion that most of what was found out during the projects *could* have been known beforehand? I agree that in retrospect it usually turns out to be quite clear what decisions should have been made, but such findings have little relevance for how decisions should be made *ex ante*.

Scherer does give some particular reasons for believing that development is not nearly as unpredictable as it seems. One of his main arguments is that there is usually a good deal of optimism in the initial estimates of performance, cost, and availability; another is that military planners, in their quest for improvements in performance, intentionally do not stick to their initial plans. However, I find neither of these arguments very convincing.

As for the first, I will certainly agree that initial estimates are almost invariably optimistic. But if there are knowledgeable military decision makers who know what particular discount factors to apply to this estimate of cost, that estimate of performance, and so forth, I would like to know who they are. Our findings indicate that initial estimates contain not only a good deal of bias but also a good deal of variance. The problem, as it usually presents itself to the decision maker, is whether a particular system is likely to cost only 50 per cent more than the initial estimate, or 2 or 3 times as much; whether the bombing accuracy of a new missile might turn out to be, if not as good as that achieved by manned aircraft, "nearly" as good; whether the technical difficulties are likely to prove modest or major. The decision maker can, of course, consult impartial experts, and frequently he does. But having read dozens of reports prepared by various expert committees, I find that the chief result of such appraisals is only to tell the decision maker that, until development has been carried forward much further, the things he would like to know will remain very conjectural matters. This is not to say that such advice is often heeded.

Now let me turn to Scherer's second argument. I do not deny that plans are sometimes changed in the course of development in order to achieve what might be regarded as incremental gains in performance. But that such a striving for perfection is a major factor in explaining the many changes that do take place during the course of development, I simply refuse to believe. The particular evidence that Scherer gives in this connection does not seem to me, at least, to be very persuasive. He says that the six fighter programs I discussed were developed in a period of relatively little international tension, and when "we were not yet worried about the threat of high performance Russian

bombers." I don't know whom he is referring to when he says "we," but it certainly is not the military officials who actually made the decisions. They regarded the development of supersonic fighters an urgent necessity and, as a matter of fact, four of the six programs were undertaken as high priority, development-production programs. Undertaking a program as a crash effort does not in itself, however, make the development problems involved any less real. The particular fighter whose development was most rapidly pressed, and whose initial funding was the largest, happened to be the one that underwent the most extensive modifications before it was put into operational use. And the modifications it underwent can hardly be described as modifications required only to get incremental gains in performance.

If it were normal in military development not to be in a hurry and to constantly strive for perfection, one would expect that occasionally the performance finally achieved would turn out to be better than that initially planned and that, when airplanes and missiles finally did go into operational use, they would more often than not, be in fine working order. Yet, we have been able to find very few cases when the first happened, and not many when the second happened. In the cases we looked into (which were many more than the fighter plane examples cited), many of the changes in the original designs were made to keep the programs from slipping more than they otherwise would have slipped to prevent major degradation in performance, and to assure that the equipment would be reasonably reliable (when I say "reasonably reliable" I do not mean the comparatively nearly perfect working order we require in our TV sets). These changes in turn often resulted in large cost increases.

When Scherer says that when a weapon is rushed into production, all that is typically sacrificed is "just a few percentage points in relevant performance values," I wonder if he includes among those relevant values: does the thing work? The experiences we have examined strongly suggest that there is much to be gained by rushing aircraft and missiles into test, but this is an entirely different matter from rushing them into production.

The basic difference between Scherer and myself comes out, I think, when he argues for the necessity of putting tight constraints on the engineers who are developing a system. There comes a point, he says, when the technical uncertainties will have been largely overcome and when detailed decisions on the final configuration have to be made. That this is so, I certainly agree, but I wonder if he is not defining

development as an activity that starts at that point—presumably all that precedes it is research.

If one defines development in this way, certainly the nature of the activity will appear to be very different from its nature according to our broader definition. My reasons for not defining development as Scherer has are, first, that such a definition does not agree with what people in the business usually have in mind when they speak about development (though he prefers to regard designing a new compressor as a research problem, programs commonly referred to as development programs often involve several problems of comparable difficulty); and, second, that if economists accept this narrower definition, I am afraid they will not be able to get very far in explaining how progress comes about. Though the tools that economists have worked out are much more applicable if development is defined as an activity consisting of assembling known pieces of technology, I hope that we will not regard that as sufficient reason to so limit the scope of our inquiry.

Finally, let me say that I think Scherer is much too complimentary when he regards our work as an attempt to formulate “a single best approach to the development of advanced systems . . .”. Actually, the aim of the paper I wrote was much less pretentious than that; it was simply to indicate that the character of the decision making problem in development is very different from that in production. This I would regard as a far cry from a positive theory and as the main shortcoming of our work thus far.

REJOINDER BY SCHERER

It is good that Burton Klein and I are still in disagreement, for disagreement is often the prelude to progress. What we need now, I think, are two things. First, we need a substantial body of empirical material on specific advanced system developments. Then it would be possible for us (and more important, for others interested in the economics of technological innovation) at least to isolate the factual and interpretational reasons for our disagreement. Unfortunately, proprietary and military security limitations are at present serious barriers to the publication of illuminating development case studies. But until more studies like those by Thomas Marschak and Richard Nelson are available, I am afraid we shall be led into the futile business of disagreeing on facts. Second, we need a better understanding of the

kinds of trade-offs that can be made between performance, reliability, cost, and development time. Let me try to clear up some misunderstanding in this latter area.

As I said in my comment, substantial time and cost savings may often be obtained by making relatively small performance sacrifices. This does not imply only foregoing increments of performance above one's original objectives. It also means making do with a little less than was originally sought. If the original goals with respect to performance are optimistic, as is so often the case, it might be better to fall short and save considerable time and money. I realize there are institutional barriers to making such performance sacrifices. However, in the short space allowed for this reply, the medicine prescribed could be nothing more than a sugar pill, so I will not attempt to write a prescription.

Now, a word about reliability, that is, does the thing work? Let me say, first, that my "ordered into production as soon as possible" and Klein's "rushing into production" connote different time scales, at least in my mind. Achieving substantial performance advances and obtaining high reliability are usually antithetical, and a good deal of development effort is required to satisfy both objectives together. If, in the middle of a development program with ambitious performance goals, the system is rushed into production, it probably will not work, at least not very well. When early production is a major objective, then it is prudent *at the outset* to limit one's performance advancement goals to what can be obtained with high certainty. This means specifying already proven concepts and components, in other words, making early configuration decisions—just what Klein apparently condemns. If advancing the state of the art is the primary objective, then one *should* keep his technical plan flexible. But if a decision maker tries to have his cake and eat it, that is, to make big performance gains and still have an operational weapon in a very short time, he is begging for disappointment unless he is willing to spend a great deal on the insurance of multiple approaches. These are the most important optimization decisions in systems development, and they must be made early in the game. As the development progresses, the opportunities for making acceptable time-performance-reliability trade-offs become more and more constrained.

I do not want to argue with Klein over definitions of development. I do, however, want to make sure we are talking about the same thing. Plainly, there are all kinds of developments in terms of an a priori

uncertainty classification. I have been referring to systems development, which might be defined very loosely as assembling known (although not always) pieces of technology. Most of Klein's examples are drawn from this sort of development, and most of our R and D money is spent on such activity. That behavior should be different for other kinds of development cannot be denied.

REJOINDER BY KLEIN

If by "known technology" Frederic Scherer means "technology that does not involve any new scientific principles," then I agree that development can be defined as an activity involving the assembly of known pieces of technology. But, as I tried to emphasize in writing my paper, there is quite a difference between knowing that some idea, say, for a communications satellite system, does not violate any known scientific principles, and knowing when it can be ready for use, how much it will cost, how well it will operate, etc. (see, for example, Marschak's paper). And in this more relevant sense I insist that most of the money spent on R and D has not been spent on "known" technology. It seems to me that simply by reading the newspaper accounts about the B-70 bomber, about the Polaris missile, about the satellite warning system, about the missile that is supposed to shoot down ballistic missiles, about the Saturn development program, etc., one might deduce that capabilities aimed for in development are more than a little different from the things we have today.