

SEVEN

The Methodology of Positive Economics

Milton Friedman

Milton Friedman (1912–2006) was born in Brooklyn, New York, and received his Ph.D. in economics from Columbia University. He taught at the University of Minnesota, and then for many years at the University of Chicago. After 1977, he was a Senior Research Fellow at the Hoover Institution in Stanford, California. Friedman is best known for his work in monetary theory and for his concern for free enterprise and individual liberty. Milton Friedman was awarded the Nobel Prize in economics in 1976. The following essay, which is reprinted in its entirety, is the most influential work on economic methodology of this century.

In his admirable book on *The Scope and Method of Political Economy* John Neville Keynes distinguishes among “a *positive science* . . . [,] a body of systematized knowledge concerning what is; a *normative* or *regulative science* . . . [,] a body of systematized knowledge discussing criteria of what ought to be . . . ; an *art* . . . [,] a system of rules for the attainment of a given end”; comments that “confusion between them is common and has been the source of many mischievous errors”; and urges the importance of “recognizing a distinct positive science of political economy.”¹

This [essay] is concerned primarily with certain methodological problems that arise in constructing the “distinct positive science” Keynes called for – in particular, the problem how to decide whether a suggested hypothesis or theory should be tentatively accepted as part of the “body of systematized

I have incorporated bodily in this article without special reference most of my brief “Comment” in *A Survey of Contemporary Economics*, Vol. II (B. F. Haley, ed.) (Chicago: Richard D. Irwin, Inc., 1952), pp. 455–7.

I am indebted to Dorothy S. Brady, Arthur F. Burns, and George J. Stigler for helpful comments and criticism.

From *Essays in Positive Economics*, by Milton Friedman. Chicago: University of Chicago Press, 1953. Copyright © 1953 by the University of Chicago. Reprinted by permission of the University of Chicago.

knowledge concerning what is.” But the confusion Keynes laments is still so rife and so much of a hindrance to the recognition that economics can be, and in part is, a positive science that it seems well to preface the main body of the paper with a few remarks about the relation between positive and normative economics.

I. The Relation between Positive and Normative Economics

Confusion between positive and normative economics is to some extent inevitable. The subject matter of economics is regarded by almost everyone as vitally important to himself and within the range of his own experience and competence; it is the source of continuous and extensive controversy and the occasion for frequent legislation. Self-proclaimed “experts” speak with many voices and can hardly all be regarded as disinterested; in any event, on questions that matter so much, “expert” opinion could hardly be accepted solely on faith even if the “experts” were nearly unanimous and clearly disinterested.² The conclusions of positive economics seem to be, and are, immediately relevant to important normative problems, to questions of what ought to be done and how any given goal can be attained. Laymen and experts alike are inevitably tempted to shape positive conclusions to fit strongly held normative preconceptions and to reject positive conclusions if their normative implications – or what are said to be their normative implications – are unpalatable.

Positive economics is in principle independent of any particular ethical position or normative judgments. As Keynes says, it deals with “what is,” not with “what ought to be.” Its task is to provide a system of generalizations that can be used to make correct predictions about the consequences of any change in circumstances. Its performance is to be judged by the precision, scope, and conformity with experience of the predictions it yields. In short, positive economics is, or can be, an “objective” science, in precisely the same sense as any of the physical sciences. Of course, the fact that economics deals with the interrelations of human beings, and that the investigator is himself part of the subject matter being investigated in a more intimate sense than in the physical sciences, raises special difficulties in achieving objectivity at the same time that it provides the social scientist with a class of data not available to the physical scientist. But neither the one nor the other is, in my view, a fundamental distinction between the two groups of sciences.³

Normative economics and the art of economics, on the other hand, cannot be independent of positive economics. Any policy conclusion necessarily rests on a prediction about the consequences of doing one thing rather

than another, a prediction that must be based – implicitly or explicitly – on positive economics. There is not, of course, a one-to-one relation between policy conclusions and the conclusions of positive economics; if there were, there would be no separate normative science. Two individuals may agree on the consequences of a particular piece of legislation. One may regard them as desirable on balance and so favor the legislation; the other, as undesirable and so oppose the legislation.

I venture the judgment, however, that currently in the Western world, and especially in the United States, differences about economic policy among disinterested citizens derive predominantly from different predictions about the economic consequences of taking action – differences that in principle can be eliminated by the progress of positive economics – rather than from fundamental differences in basic values, differences about which men can ultimately only fight. An obvious and not unimportant example is minimum-wage legislation. Underneath the welter of arguments offered for and against such legislation there is an underlying consensus on the objective of achieving a “living wage” for all, to use the ambiguous phrase so common in such discussions. The difference of opinion is largely grounded on an implicit or explicit difference in predictions about the efficacy of this particular means in furthering the agreed-on end. Proponents believe (predict) that legal minimum wages diminish poverty by raising the wages of those receiving less than the minimum wage as well as of some receiving more than the minimum wage without any counterbalancing increase in the number of people entirely unemployed or employed less advantageously than they otherwise would be. Opponents believe (predict) that legal minimum wages increase poverty by increasing the number of people who are unemployed or employed less advantageously and that this more than offsets any favorable effect on the wages of those who remain employed. Agreement about the economic consequences of the legislation might not produce complete agreement about its desirability, for differences might still remain about its political or social consequences; but, given agreement on objectives, it would certainly go a long way toward producing consensus.

Closely related differences in positive analysis underlie divergent views about the appropriate role and place of trade-unions and the desirability of direct price and wage controls and of tariffs. Different predictions about the importance of so-called “economics of scale” account very largely for divergent views about the desirability or necessity of detailed government regulation of industry and even of socialism rather than private enterprise. And this list could be extended indefinitely.⁴ Of course, my judgment that the major differences about economic policy in the Western world are of

this kind is itself a “positive” statement to be accepted or rejected on the basis of empirical evidence.

If this judgment is valid, it means that a consensus on “correct” economic policy depends much less on the progress of normative economics proper than on the progress of a positive economics yielding conclusions that are, and deserve to be, widely accepted. It means also that a major reason for distinguishing positive economics sharply from normative economics is precisely the contribution that can thereby be made to agreement about policy.

II. Positive Economics

The ultimate goal of a positive science is the development of a “theory” or “hypothesis” that yields valid and meaningful (i.e., not truistic) predictions about phenomena not yet observed. Such a theory is, in general, a complex intermixture of two elements. In part, it is a “language” designed to promote “systematic and organized methods of reasoning.”⁵ In part, it is a body of substantive hypotheses designed to abstract essential features of complex reality.

Viewed as a language, theory has no substantive content; it is a set of tautologies. Its function is to serve as a filing system for organizing empirical material and facilitating our understanding of it; and the criteria by which it is to be judged are those appropriate to a filing system. Are the categories clearly and precisely defined? Are they exhaustive? Do we know where to file each individual item, or is there considerable ambiguity? Is the system of headings and subheadings so designed that we can quickly find an item we want, or must we hunt from place to place? Are the items we shall want to consider jointly filed together? Does the filing system avoid elaborate cross-references?

The answers to these questions depend partly on logical, partly on factual, considerations. The canons of formal logic alone can show whether a particular language is complete and consistent, that is, whether propositions in the language are “right” or “wrong.” Factual evidence alone can show whether the categories of the “analytical filing system” have a meaningful empirical counterpart, that is, whether they are useful in analyzing a particular class of concrete problems.⁶ The simple example of “supply” and “demand” illustrates both this point and the preceding list of analogical questions. Viewed as elements of the language of economic theory, these are the two major categories into which factors affecting the relative prices of products or factors of production are classified. The usefulness of the

dichotomy depends on the “empirical generalization that an enumeration of the forces affecting demand in any problem and of the forces affecting supply will yield two lists that contain few items in common.”⁷ Now this generalization is valid for markets like the final market for a consumer good. In such a market there is a clear and sharp distinction between the economic units that can be regarded as demanding the product and those that can be regarded as supplying it. There is seldom much doubt whether a particular factor should be classified as affecting supply, on the one hand, or demand, on the other; and there is seldom much necessity for considering cross-effects (cross-references) between the two categories. In these cases the simple and even obvious step of filing the relevant factors under the headings of “supply” and “demand” effects a great simplification of the problem and is an effective safeguard against fallacies that otherwise tend to occur. But the generalization is not always valid. For example, it is not valid for the day-to-day fluctuations of prices in a primarily speculative market. Is a rumor of an increased excess-profits tax, for example, to be regarded as a factor operating primarily on today’s supply of corporate equities in the stock market or on today’s demand for them? In similar fashion, almost every factor can with about as much justification be classified under the heading “supply” as under the heading “demand.” These concepts can still be used and may not be entirely pointless; they are still “right” but clearly less useful than in the first example because they have no meaningful empirical counterpart.

Viewed as a body of substantive hypotheses, theory is to be judged by its predictive power for the class of phenomena which it is intended to “explain.” Only factual evidence can show whether it is “right” or “wrong” or, better, tentatively “accepted” as valid or “rejected.” As I shall argue at greater length below, the only relevant test of the *validity* of a hypothesis is comparison of its predictions with experience. The hypothesis is rejected if its predictions are contradicted (“frequently” or more often than predictions from an alternative hypothesis); it is accepted if its predictions are not contradicted; great confidence is attached to it if it has survived many opportunities for contradiction. Factual evidence can never “prove” a hypothesis; it can only fail to disprove it, which is what we generally mean when we say, somewhat inexactly, that the hypothesis has been “confirmed” by experience.

To avoid confusion, it should perhaps be noted explicitly that the “predictions” by which the validity of a hypothesis is tested need not be about phenomena that have not yet occurred, that is, need not be forecasts of future events; they may be about phenomena that have occurred but observations on which have not yet been made or are not known to the person making

the prediction. For example, a hypothesis may imply that such and such must have happened in 1906, given some other known circumstances. If a search of the records reveals that such and such did happen, the prediction is confirmed; if it reveals that such and such did not happen, the prediction is contradicted.

The validity of a hypothesis in this sense is not by itself a sufficient criterion for choosing among alternative hypotheses. Observed facts are necessarily finite in number; possible hypotheses, infinite. If there is one hypothesis that is consistent with the available evidence, there are always an infinite number that are.⁸ For example, suppose a specific excise tax on a particular commodity produces a rise in price equal to the amount of the tax. This is consistent with competitive conditions, a stable demand curve, and a horizontal and stable supply curve. But it is also consistent with competitive conditions and a positively or negatively sloping supply curve with the required compensating shift in the demand curve or the supply curve; with monopolistic conditions, constant marginal costs, and stable demand curve, of the particular shape required to produce this result; and so on indefinitely. Additional evidence with which the hypothesis is to be consistent may rule out some of these possibilities; it can never reduce them to a single possibility alone capable of being consistent with the finite evidence. The choice among alternative hypotheses equally consistent with the available evidence must to some extent be arbitrary, though there is general agreement that relevant considerations are suggested by the criteria “simplicity” and “fruitfulness,” themselves notions that defy completely objective specification. A theory is “simpler” the less the initial knowledge needed to make a prediction within a given field of phenomena; it is more “fruitful” the more precise the resulting prediction, the wider the area within which the theory yields predictions, and the more additional lines for further research it suggests. Logical completeness and consistency are relevant but play a subsidiary role; their function is to assure that the hypothesis says what it is intended to say and does so alike for all users – they play the same role here as checks for arithmetical accuracy do in statistical computations.

Unfortunately, we can seldom test particular predictions in the social sciences by experiments explicitly designed to eliminate what are judged to be the most important disturbing influences. Generally, we must rely on evidence cast up by the “experiments” that happen to occur. The inability to conduct so-called “controlled experiments” does not, in my view, reflect a basic difference between the social and physical sciences both because it is not peculiar to the social sciences – witness astronomy – and because the distinction between a controlled experiment and uncontrolled experience is

at best one of degree. No experiment can be completely controlled, and every experience is partly controlled, in the sense that some disturbing influences are relatively constant in the course of it.

Evidence cast up by experience is abundant and frequently as conclusive as that from contrived experiments; thus the inability to conduct experiments is not a fundamental obstacle to testing hypotheses by the success of their predictions. But such evidence is far more difficult to interpret. It is frequently complex and always indirect and incomplete. Its collection is often arduous, and its interpretation generally requires subtle analysis and involved chains of reasoning, which seldom carry real conviction. The denial to economics of the dramatic and direct evidence of the "crucial" experiment does hinder the adequate testing of hypotheses; but this is much less significant than the difficulty it places in the way of achieving a reasonably prompt and wide consensus on the conclusions justified by the available evidence. It renders the weeding-out of unsuccessful hypotheses slow and difficult. They are seldom downed for good and are always cropping up again.

There is, of course, considerable variation in these respects. Occasionally, experience casts up evidence that is about as direct, dramatic, and convincing as any that could be provided by controlled experiments. Perhaps the most obviously important example is the evidence from inflations on the hypothesis that a substantial increase in the quantity of money within a relatively short period is accompanied by a substantial increase in prices. Here the evidence is dramatic, and the chain of reasoning required to interpret it is relatively short. Yet, despite numerous instances of substantial rises in prices, their essentially one-to-one correspondence with substantial rises in the stock of money, and the wide variation in other circumstances that might appear to be relevant, each new experience of inflation brings forth vigorous contentions, and not only by the lay public, that the rise in the stock of money is either an incidental effect of a rise in prices produced by other factors or a purely fortuitous and unnecessary concomitant of the price rise.

One effect of the difficulty of testing substantive economic hypotheses has been to foster a retreat into purely formal or tautological analysis.⁹ As already noted, tautologies have an extremely important place in economics and other sciences as a specialized language or "analytical filing system." Beyond this, formal logic and mathematics, which are both tautologies, are essential aids in checking the correctness of reasoning, discovering the implications of hypotheses, and determining whether supposedly different hypotheses may not really be equivalent or wherein the differences lie.

But economic theory must be more than a structure of tautologies if it is to be able to predict and not merely describe the consequences of action; if it is to be something different from disguised mathematics.¹⁰ And the usefulness of the tautologies themselves ultimately depends, as noted above, on the acceptability of the substantive hypotheses that suggest the particular categories into which they organize the refractory empirical phenomena.

A more serious effect of the difficulty of testing economic hypotheses by their predictions is to foster misunderstanding of the role of empirical evidence in theoretical work. Empirical evidence is vital at two different, though closely related, stages: in constructing hypotheses and in testing their validity. Full and comprehensive evidence on the phenomena to be generalized or “explained” by a hypothesis, besides its obvious value in suggesting new hypotheses, is needed to assure that a hypothesis explains what it sets out to explain – that its implications for such phenomena are not contradicted in advance by experience that has already been observed.¹¹ Given that the hypothesis is consistent with the evidence at hand, its further testing involves deducing from it new facts capable of being observed but not previously known and checking these deduced facts against additional empirical evidence. For this test to be relevant, the deduced facts must be about the class of phenomena the hypothesis is designed to explain; and they must be well enough defined so that observation can show them to be wrong.

The two stages of constructing hypotheses and testing their validity are related in two different respects. In the first place, the particular facts that enter at each stage are partly an accident of the collection of data and the knowledge of the particular investigator. The facts that serve as a test of the implications of a hypothesis might equally well have been among the raw material used to construct it, and conversely. In the second place, the process never begins from scratch; the so-called “initial stage” itself always involves comparison of the implications of an earlier set of hypotheses with observation; the contradiction of these implications is the stimulus to the construction of new hypotheses or revision of old ones. So the two methodologically distinct stages are always proceeding jointly.

Misunderstanding about this apparently straightforward process centers on the phrase “the class of phenomena the hypothesis is designed to explain.” The difficulty in the social sciences of getting new evidence for this class of phenomena and of judging its conformity with the implications of the hypothesis makes it tempting to suppose that other, more readily available, evidence is equally relevant to the validity of the hypothesis – to suppose that hypotheses have not only “implications” but also “assumptions” and that

the conformity of these “assumptions” to “reality” is a test of the validity of the hypothesis *different from* or *additional to* the test by implications, this widely held view is fundamentally wrong and productive of much mischief. Far from providing an easier means for sifting valid from invalid hypotheses, it only confuses the issue, promotes misunderstanding about the significance of empirical evidence for economic theory, produces a misdirection of much intellectual effort devoted to the development of positive economics, and impedes the attainment of consensus on tentative hypotheses in positive economics.

In so far as a theory can be said to have “assumptions” at all, and in so far as their “realism” can be judged independently of the validity of predictions, the relation between the significance of a theory and the “realism” of its “assumptions” is almost the opposite of that suggested by the view under criticism. Truly important and significant hypotheses will be found to have “assumptions” that are wildly inaccurate descriptive representations of reality, and, in general, the more significant the theory, the more unrealistic the assumptions (in this sense).¹² The reason is simple. A hypothesis is important if it “explains” much by little, that is, if it abstracts the common and crucial elements from the mass of complex and detailed circumstances surrounding the phenomena to be explained and permits valid predictions on the basis of them alone. To be important, therefore, a hypothesis must be descriptively false in its assumptions; it takes account of, and accounts for, none of the many other attendant circumstances, since its very success shows them to be irrelevant for the phenomena to be explained.

To put this point less paradoxically, the relevant question to ask about the “assumptions” of a theory is not whether they are descriptively “realistic,” for they never are, but whether they are sufficiently good approximations for the purpose in hand. And this question can be answered only by seeing whether the theory works, which means whether it yields sufficiently accurate predictions. The two supposedly independent tests thus reduce to one test.

The theory of monopolistic and imperfect competition is one example of the neglect in economic theory of these propositions. The development of this analysis was explicitly motivated, and its wide acceptance and approval largely explained, by the belief that the assumptions of “perfect competition” or “perfect monopoly” said to underlie neoclassical economic theory are a false image of reality. And this belief was itself based almost entirely on the directly perceived descriptive inaccuracy of the assumptions rather than on any recognized contradiction of predictions derived from neoclassical economic theory. The lengthy discussion on marginal analysis in the

American Economic Review some years ago is an even clearer, though much less important, example. The articles on both sides of the controversy largely neglect what seems to me clearly the main issue – the conformity to experience of the implications of the marginal analysis – and concentrate on the largely irrelevant question whether businessmen do or do not in fact reach their decisions by consulting schedules, or curves, or multivariable functions showing marginal cost and marginal revenue.¹³ Perhaps these two examples, and the many others they readily suggest, will serve to justify a more extensive discussion of the methodological principles involved than might otherwise seem appropriate.

III. Can a Hypothesis be Tested by the Realism of its Assumptions?

We may start with a simple physical example, the law of falling bodies. It is an accepted hypothesis that the acceleration of a body dropped in a vacuum is a constant – g , or approximately 32 feet per second per second on the earth – and is independent of the shape of the body, the manner of dropping it, etc. This implies that the distance traveled by a falling body in any specified time is given by the formula $s = 1/2 gt^2$, where s is the distance traveled in feet and t is time in seconds. The application of this formula to a compact ball dropped from the roof of a building is equivalent to saying that a ball so dropped behaves *as if* it were falling in a vacuum. Testing this hypothesis by its assumptions presumably means measuring the actual air pressure and deciding whether it is close enough to zero. At sea level the air pressure is about 15 pounds per square inch. Is 15 sufficiently close to zero for the difference to be judged insignificant? Apparently it is, since the actual time taken by a compact ball to fall from the roof of a building to the ground is very close to the time given by the formula. Suppose, however, that a feather is dropped instead of a compact ball. The formula then gives wildly inaccurate results. Apparently, 15 pounds per square inch is significantly different from zero for a feather but not for a ball. Or, again, suppose the formula is applied to a ball dropped from an airplane at an altitude of 30,000 feet. The air pressure at this altitude is decidedly less than 15 pounds per square inch. Yet, the actual time of fall from 30,000 feet to 20,000 feet, at which point the air pressure is still much less than at sea level, will differ noticeably from the time predicted by the formula – much more noticeably than the time taken by a compact ball to fall from the roof of a building to the ground. According to the formula, the velocity of the ball should be gt and should therefore increase steadily. In fact, a ball dropped at 30,000 feet will reach its top velocity well before it hits the ground. And similarly with other implications of the formula.

The initial question whether 15 is sufficiently close to zero for the difference to be judged insignificant is clearly a foolish question by itself. Fifteen pounds per square inch is 2,160 pounds per square foot, or 0.0075 ton per square inch. There is no possible basis for calling these numbers “small” or “large” without some external standard of comparison. And the only relevant standard of comparison is the air pressure for which the formula does or does not work under a given set of circumstances. But this raises the same problem at a second level. What is the meaning of “does or does not work”? Even if we could eliminate errors of measurement, the measured time of fall would seldom if ever be precisely equal to the computed time of fall. How large must the difference between the two be to justify saying that the theory “does not work”? Here there are two important external standards of comparison. One is the accuracy achievable by an alternative theory with which this theory is being compared and which is equally acceptable on all other grounds. The other arises when there exists a theory that is known to yield better predictions but only at a greater cost. The gains from greater accuracy, which depend on the purpose in mind, must then be balanced against the costs of achieving it.

The example illustrates both the impossibility of testing a theory by its assumptions and also the ambiguity of the concept “the assumptions of a theory.” The formula $s = 1/2 gt^2$ is valid for bodies falling in a vacuum and can be derived by analyzing the behavior of such bodies. It can therefore be stated: under a wide range of circumstances, bodies that fall in the actual atmosphere behave *as if* they were falling in a vacuum. In the language so common in economics this would be rapidly translated into: the formula assumes a vacuum. Yet it clearly does no such thing. What it does say is that in many cases the existence of air pressure, the shape of the body, the name of the person dropping the body, the kind of mechanism used to drop the body, and a host of other attendant circumstances have no appreciable effect on the distance the body falls in a specified time. The hypothesis can readily be rephrased to omit all mention of a vacuum: under a wide range of circumstances, the distance a body falls in a specified time is given by the formula $s = 1/2 gt^2$. The history of this formula and its associated physical theory aside, is it meaningful to say that it assumes a vacuum? For all I know there may be other sets of assumptions that would yield the same formula. The formula is accepted because it works, not because we live in an approximate vacuum – whatever that means.

The important problem in connection with the hypothesis is to specify the circumstances under which the formula works or, more precisely, the general magnitude of the error in its predictions under various circumstances. Indeed, as is implicit in the above rephrasing of the hypothesis, such

a specification is not one thing and the hypothesis another. The specification is itself an essential part of the hypothesis, and it is a part that is peculiarly likely to be revised and extended as experience accumulates.

In the particular case of falling bodies a more general, though still incomplete, theory is available, largely as a result of attempts to explain the errors of the simple theory, from which the influence of some of the possible disturbing factors can be calculated and of which the simple theory is a special case. However, it does not always pay to use the more general theory because the extra accuracy it yields may not justify the extra cost of using it, so the question under what circumstances the simpler theory works “well enough” remains important. Air pressure is one, but only one, of the variables that define these circumstances; the shape of the body, the velocity attained, and still other variables are relevant as well. One way of interpreting the variables other than air pressure is to regard them as determining whether a particular departure from the “assumption” of a vacuum is or is not significant. For example, the difference in shape of the body can be said to make 15 pounds per square inch significantly different from zero for a feather but not for a compact ball dropped a moderate distance. Such a statement must, however, be sharply distinguished from the very different statement that the theory does not work for a feather because its assumptions are false. The relevant relation runs the other way: the assumptions are false for a feather because the theory does not work. This point needs emphasis, because the entirely valid use of “assumptions” in *specifying* the circumstances for which theory holds is frequently, and erroneously, interpreted to mean that the assumptions can be used to *determine* the circumstances for which a theory holds, and has, in this way, been an important source of the belief that a theory can be tested by its assumptions.

Let us turn now to another example, this time a constructed one designed to be an analogue of many hypotheses in the social sciences. Consider the density of leaves around a tree. I suggest the hypothesis that the leaves are positioned as if each leaf deliberately sought to maximize the amount of sunlight it receives, given the position of its neighbors, as if it knew the physical laws determining the amount of sunlight that would be received in various positions and could move rapidly or instantaneously from any one position to any other desired and unoccupied position.¹⁴ Now some of the more obvious implications of this hypothesis are clearly consistent with experience: for example, leaves are in general denser on the south than on the north side of trees but, as the hypothesis implies, less so or not at all on the northern slope of a hill or when the south side of the trees is shaded in some other way. Is the hypothesis rendered unacceptable or

invalid because, so far as we know, leaves do not “deliberate” or consciously “seek,” have not been to school and learned the relevant laws of science or the mathematics required to calculate the “optimum” position, and cannot move from position to position? Clearly, none of these contradictions of the hypothesis is vitally relevant; the phenomena involved are not within the “class of phenomena the hypothesis is designed to explain”; the hypothesis does not assert that leaves do these things but only that their density is the same *as if* they did. Despite the apparent falsity of the “assumptions” of the hypothesis, it has great plausibility because of the conformity of its implications with observation. We are inclined to “explain” its validity on the ground that sunlight contributes to the growth of leaves and that hence leaves will grow denser or more putative leaves survive where there is more sun, so the result achieved by purely passive adaptation to external circumstances is the same as the result that would be achieved by deliberate accommodation to them. This alternative hypothesis is more attractive than the constructed hypothesis not because its “assumptions” are more “realistic” but rather because it is part of a more general theory that applies to a wider variety of phenomena, of which the position of leaves around a tree is a special case, has more implications capable of being contradicted, and has failed to be contradicted under a wider variety of circumstances. The direct evidence for the growth of leaves is in this way strengthened by the indirect evidence from the other phenomena to which the more general theory applies.

The constructed hypothesis is presumably valid, that is, yields “sufficiently” accurate predictions about the density of leaves, only for a particular class of circumstances. I do not know what these circumstances are or how to define them. It seems obvious, however, that in this example the “assumptions” of the theory will play no part in specifying them: the kind of tree, the character of the soil, etc., are the types of variables that are likely to define its range of validity, not the ability of the leaves to do complicated mathematics or to move from place to place.

A largely parallel example involving human behavior has been used elsewhere by Savage and me.¹⁵ Consider the problem of predicting the shots made by an expert billiard player. It seems not at all unreasonable that excellent predictions would be yielded by the hypothesis that the billiard player made his shots *as if* he knew the complicated mathematical formulas that would give the optimum directions of travel, could estimate accurately by eye the angles, etc., describing the location of the balls, could make lightning calculations from the formulas, and could then make the balls travel in the direction indicated by the formulas. Our confidence in this hypothesis is not based on the belief that billiard players, even expert ones, can or do go

through the process described; it derives rather from the belief that, unless in some way or other they were capable of reaching essentially the same result, they would not in fact be *expert* billiard players.

It is only a short step from these examples to the economic hypothesis that under a wide range of circumstances individual firms behave *as if* they were seeking rationally to maximize their expected returns (generally if misleadingly called “profits”)¹⁶ and had full knowledge of the data needed to succeed in this attempt; *as if*, that is, they knew the relevant cost and demand functions, calculated marginal cost and marginal revenue from all actions open to them, and pushed each line of action to the point at which the relevant marginal cost and marginal revenue were equal. Now, of course, businessmen do not actually and literally solve the system of simultaneous equations in terms of which the mathematical economist finds it convenient to express this hypothesis, any more than leaves or billiard players explicitly go through complicated mathematical calculations or falling bodies decide to create a vacuum. The billiard player, if asked how he decides where to hit the ball, may say that he “just figures it out” but then also rubs a rabbit’s foot just to make sure; and the businessman may well say that he prices at average cost, with of course some minor deviations when the market makes it necessary. The one statement is about as helpful as the other, and neither is a relevant test of the associated hypothesis.

Confidence in the maximization-of-returns hypothesis is justified by evidence of a very different character. This evidence is in part similar to that adduced on behalf of the billiard-player hypothesis – unless the behavior of businessmen in some way or other approximated behavior consistent with the maximization of returns, it seems unlikely that they would remain in business for long. Let the apparent immediate determinant of business behavior be anything at all – habitual reaction, random chance, or what-not. Whenever this determinant happens to lead to behavior consistent with rational and informed maximization of returns, the business will prosper and acquire resources with which to expand; whenever it does not, the business will tend to lose resources and can be kept in existence only by the addition of resources from outside. The process of “natural selection” thus helps to validate the hypothesis – or, rather, given natural selection, acceptance of the hypothesis can be based largely on the judgment that it summarizes appropriately the conditions for survival.

An even more important body of evidence for the maximization-of-returns hypothesis is experience from countless applications of the hypothesis to specific problems and the repeated failure of its implications to be contradicted. This evidence is extremely hard to document; it is scattered in

numerous memorandums, articles, and monographs concerned primarily with specific concrete problems rather than with submitting the hypothesis to test. Yet the continued use and acceptance of the hypothesis over a long period, and the failure of any coherent, self-consistent alternative to be developed and be widely accepted, is strong indirect testimony to its worth. The evidence *for* a hypothesis always consists of its repeated failure to be contradicted, continues to accumulate so long as the hypothesis is used, and by its very nature is difficult to document at all comprehensively. It tends to become part of the tradition and folklore of a science revealed in the tenacity with which hypotheses are held rather than in any textbook list of instances in which the hypothesis has failed to be contradicted.

IV. The Significance and Role of the “Assumptions” of a Theory

Up to this point our conclusions about the significance of the “assumptions” of a theory have been almost entirely negative: we have seen that a theory cannot be tested by the “realism” of its “assumptions” and that the very concept of the “assumptions” of a theory is surrounded with ambiguity. But, if this were all there is to it, it would be hard to explain the extensive use of the concept and the strong tendency that we all have to speak of the assumptions of a theory and to compare the assumptions of alternative theories. There is too much smoke for there to be no fire.

In methodology, as in positive science, negative statements can generally be made with greater confidence than positive statements, so I have less confidence in the following remarks on the significance and role of “assumptions” than in the preceding remarks. So far as I can see, the “assumptions of a theory” play three different, though related, positive roles: (a) they are often an economical mode of describing or presenting a theory; (b) they sometimes facilitate an indirect test of the hypothesis by its implications; and (c), as already noted, they are sometimes a convenient means of specifying the conditions under which the theory is expected to be valid. The first two require more extensive discussion.

A. The Use of “Assumptions” in Stating a Theory

The example of the leaves illustrates the first role of assumptions. Instead of saying that leaves seek to maximize the sunlight they receive, we could state the equivalent hypothesis, without any apparent assumptions, in the form of a list of rules for predicting the density of leaves: If a tree stands in a level field with no other trees or other bodies obstructing the rays of the

sun, then the density of leaves will tend to be such and such; if a tree is on the northern slope of a hill in the midst of a forest of similar trees, then . . . ; etc. This is clearly a far less economical presentation of the hypothesis than the statement that leaves seek to maximize the sunlight each receives. The latter statement is, in effect, a simple summary of the rules in the above list, even if the list were indefinitely extended, since it indicates both how to determine the features of the environment that are important for the particular problem and how to evaluate their effects. It is more compact and at the same time no less comprehensive.

More generally, a hypothesis or theory consists of an assertion that certain forces are, and by implication others are not, important for a particular class of phenomena and a specification of the manner of action of the forces it asserts to be important. We can regard the hypothesis as consisting of two parts: first, a conceptual world or abstract model simpler than the “real world” and containing only the forces that the hypothesis asserts to be important; second, a set of rules defining the class of phenomena for which the “model” can be taken to be an adequate representation of the “real world” and specifying the correspondence between the variables or entities in the model and observable phenomena.

These two parts are very different in character. The model is abstract and complete; it is an “algebra” or “logic.” Mathematics and formal logic come into their own in checking its consistency and completeness and exploring its implications. There is no place in the model for, and no function to be served by, vagueness, maybe’s, or approximations. The air pressure is zero, not “small,” for a vacuum; the demand curve for the product of a competitive producer is horizontal (has a slope of zero), not “almost horizontal.”

The rules for using the model, on the other hand, cannot possibly be abstract and complete. They must be concrete and in consequence incomplete – completeness is possible only in a conceptual world, not in the “real world,” however that may be interpreted. The model is the logical embodiment of the half-truth, “There is nothing new under the sun”; the rules for applying it cannot neglect the equally significant half-truth, “History never repeats itself.” To a considerable extent the rules can be formulated explicitly – most easily, though even then not completely, when the theory is part of an explicit more general theory as in the example of the vacuum theory for falling bodies. In seeking to make a science as “objective” as possible, our aim should be to formulate the rules explicitly in so far as possible and continually to widen the range of phenomena for which it is possible to do so. But, no matter how successful we may be in this attempt, there

inevitably will remain room for judgment in applying the rules. Each occurrence has some features peculiarly its own, not covered by the explicit rules. The capacity to judge that these are or are not to be disregarded, that they should or should not affect what observable phenomena are to be identified with what entities in the model, is something that cannot be taught; it can be learned but only by experience and exposure in the “right” scientific atmosphere, not by rote. It is at this point that the “amateur” is separated from the “professional” in all sciences and that the thin line is drawn which distinguishes the “crackpot” from the scientist.

A simple example may perhaps clarify this point. Euclidean geometry is an abstract model, logically complete and consistent. Its entities are precisely defined – a line is not a geometrical figure “much” longer than it is wide or deep; it is a figure whose width and depth are zero. It is also obviously “unrealistic.” There are no such things in “reality” as Euclidean points or lines or surfaces. Let us apply this abstract model to a mark made on a blackboard by a piece of chalk. Is the mark to be identified with a Euclidean line, a Euclidean surface, or a Euclidean solid? Clearly, it can appropriately be identified with a line if it is being used to represent, say, a demand curve. But it cannot be so identified if it is being used to color, say, countries on a map, for that would imply that the map would never be colored; for this purpose, the same mark must be identified with a surface. But it cannot be so identified by a manufacturer of chalk, for that would imply that no chalk would ever be used up; for his purposes, the same mark must be identified with a volume. In this simple example these judgments will command general agreement. Yet it seems obvious that, while general considerations can be formulated to guide such judgments, they can never be comprehensive and cover every possible instance; they cannot have the self-contained coherent character of Euclidean geometry itself.

In speaking of the “crucial assumptions” of a theory, we are, I believe, trying to state the key elements of the abstract model. There are generally many different ways of describing the model completely – many different sets of “postulates” which both imply and are implied by the model as a whole. These are all logically equivalent: what are regarded as axioms or postulates of a model from one point of view can be regarded as theorems from another, and conversely. The particular “assumptions” termed “crucial” are selected on grounds of their convenience in some such respects as simplicity or economy in describing the model, intuitive plausibility, or capacity to suggest, if only by implication, some of the considerations that are relevant in judging or applying the model.

B. The Use of “Assumptions” as an Indirect Test of Theory

In presenting any hypothesis, it generally seems obvious which of the series of statements used to expound it refer to assumptions and which to implications; yet this distinction is not easy to define rigorously. It is not, I believe, a characteristic of the hypothesis as such but rather of the use to which the hypothesis is to be put. If this is so, the ease of classifying statements must reflect unambiguousness in the purpose the hypothesis is designed to serve. The possibility of interchanging theorems and axioms in an abstract model implies the possibility of interchanging “implications” and “assumptions” in the substantive hypothesis corresponding to the abstract model, which is not to say that any implication can be interchanged with any assumption but only that there may be more than one set of statements that imply the rest.

For example, consider a particular proposition in the theory of oligopolistic behavior. If we assume (*a*) that entrepreneurs seek to maximize their returns by any means including acquiring or extending monopoly power, this will imply (*b*) that, when demand for a “product” is geographically unstable, transportation costs are significant, explicit price agreements illegal, and the number of producers of the product relatively small, they will tend to establish basing-point pricing systems.¹⁷ The assertion (*a*) is regarded as an assumption and (*b*) as an implication because we accept the prediction of market behavior as the purpose of the analysis. We shall regard the assumption as acceptable if we find that the conditions specified in (*b*) are generally associated with basing-point pricing, and conversely. Let us now change our purpose to deciding what cases to prosecute under the Sherman Antitrust Law’s prohibition of a “conspiracy in restraint of trade.” If we now assume (*c*) that basing-point pricing is a deliberate construction to facilitate collusion under the conditions specified in (*b*), this will imply (*d*) that entrepreneurs who participate in basing-point pricing are engaged in a “conspiracy in restraint of trade.” What was formerly an assumption now becomes an implication, and conversely. We shall now regard the assumption (*c*) as valid if we find that, when entrepreneurs participate in basing-point pricing, there generally tends to be other evidence, in the form of letters, memorandums, or the like, of what courts regard as a “conspiracy in restraint of trade.”

Suppose the hypothesis works for the first purpose, namely, the prediction of market behavior. It clearly does not follow that it will work for the second purpose, namely, predicting whether there is enough evidence of a “conspiracy in restraint of trade” to justify court action. And, conversely, if

it works for the second purpose, it does not follow that it will work for the first. Yet, in the absence of other evidence, the success of the hypothesis for one purpose – in explaining one class of phenomena – will give us greater confidence than we would otherwise have that it may succeed for another purpose – in explaining another class of phenomena. It is much harder to say how much greater confidence it justifies. For this depends on how closely related we judge the two classes of phenomena to be, which itself depends in a complex way on similar kinds of indirect evidence, that is, on our experience in other connections in explaining by single theories phenomena that are in some sense similarly diverse.

To state the point more generally, what are called the assumptions of a hypothesis can be used to get some indirect evidence on the acceptability of the hypothesis in so far as the assumptions can themselves be regarded as implications of the hypothesis, and hence their conformity with reality as a failure of some implications to be contradicted, or in so far as the assumptions may call to mind other implications of the hypothesis susceptible to casual empirical observation.¹⁸ The reason this evidence is indirect is that the assumptions or associated implications generally refer to a class of phenomena different from the class which the hypothesis is designed to explain; indeed, as is implied above, this seems to be the chief criterion we use in deciding which statements to term “assumptions” and which to term “implications.” The weight attached to this indirect evidence depends on how closely related we judge the two classes of phenomena to be.

Another way in which the “assumptions” of a hypothesis can facilitate its indirect testing is by bringing out its kinship with other hypotheses and thereby making the evidence on their validity relevant to the validity of the hypothesis in question. For example, a hypothesis is formulated for a particular class of behavior. This hypothesis can, as usual, be stated without specifying any “assumptions.” But suppose it can be shown that it is equivalent to a set of assumptions including the assumption that man seeks his own interest. The hypothesis then gains indirect plausibility from the success for other classes of phenomena of hypotheses that can also be said to make this assumption; at least, what is being done here is not completely unprecedented or unsuccessful in all other uses. In effect, the statement of assumptions so as to bring out a relationship between superficially different hypotheses is a step in the direction of a more general hypothesis.

This kind of indirect evidence from related hypotheses explains in large measure the difference in the confidence attached to a particular hypothesis by people with different backgrounds. Consider, for example, the hypothesis that the extent of racial or religious discrimination in employment in

a particular area or industry is closely related to the degree of monopoly in the industry or area in question; that, if the industry is competitive, discrimination will be significant only if the race or religion of employees affects either the willingness of other employees to work with them or the acceptability of the product to customers and will be uncorrelated with the prejudices of employers.¹⁹ This hypothesis is far more likely to appeal to an economist than to a sociologist. It can be said to “assume” single-minded pursuit of pecuniary self-interest by employers in competitive industries; and this “assumption” works well in a wide variety of hypotheses in economics bearing on many of the mass phenomena with which economics deals. It is therefore likely to seem reasonable to the economist that it may work in this case as well. On the other hand, the hypotheses to which the sociologist is accustomed have a very different kind of model or ideal world, in which singleminded pursuit of pecuniary self-interest plays a much less important role. The indirect evidence available to the sociologist on this hypothesis is much less favorable to it than the indirect evidence available to the economist; he is therefore likely to view it with greater suspicion.

Of course, neither the evidence of the economist nor that of the sociologist is conclusive. The decisive test is whether the hypothesis works for the phenomena it purports to explain. But a judgment may be required before any satisfactory test of this kind has been made, and, perhaps, when it cannot be made in the near future, in which case, the judgment will have to be based on the inadequate evidence available. In addition, even when such a test can be made, the background of the scientists is not irrelevant to the judgments they reach. There is never certainty in science, and the weight of evidence for or against a hypothesis can never be assessed completely “objectively.” The economist will be more tolerant than the sociologist in judging conformity of the implications of the hypothesis with experience, and he will be persuaded to accept the hypothesis tentatively by fewer instances of “conformity.”

V. Some Implications for Economic Issues

The abstract methodological issues we have been discussing have a direct bearing on the perennial criticism of “orthodox” economic theory as “unrealistic” as well as on the attempts that have been made to reformulate theory to meet this charge. Economics is a “dismal” science because it assumes man to be selfish and money-grubbing, “a lightning calculator of pleasures and pains, who oscillates like a homogeneous globule of desire of happiness under the impulse of stimuli that shift him about the area, but leave him

intact”;²⁰ it rests on outmoded psychology and must be reconstructed in line with each new development in psychology; it assumes men, or at least businessmen, to be “in a continuous state of ‘alert,’ ready to change prices and/or pricing rules whenever their sensitive intuitions . . . detect a change in demand and supply conditions”;²¹ it assumes markets to be perfect, competition to be pure, and commodities, labor, and capital to be homogeneous.

As we have seen, criticism of this type is largely beside the point unless supplemented by evidence that a hypothesis differing in one or another of these respects from the theory being criticized yields better predictions for as wide a range of phenomena. Yet most such criticism is not so supplemented; it is based almost entirely on supposedly directly perceived discrepancies between the “assumptions” and the “real world.” A particularly clear example is furnished by the recent criticisms of the maximization-of-returns hypothesis on the grounds that businessmen do not and indeed cannot behave as the theory “assumes” they do. The evidence cited to support this assertion is generally taken either from the answers given by businessmen to questions about the factors affecting their decisions – a procedure for testing economic theories that is about on a par with testing theories of longevity by asking octogenarians how they account for their long life – or from descriptive studies of the decision-making activities of individual firms.²² Little if any evidence is ever cited on the conformity of businessmen’s actual market behavior – what they do rather than what they say they do – with the implications of the hypothesis being criticized, on the one hand, and an alternative hypothesis, on the other.

A theory or its “assumptions” cannot possibly be thoroughly “realistic” in the immediate descriptive sense so often assigned to this term. A completely “realistic” theory of the wheat market would have to include not only the conditions directly underlying the supply and demand for wheat but also the kind of coins or credit instruments used to make exchanges; the personal characteristics of wheat-traders such as the color of each trader’s hair and eyes, his antecedents and education, the number of members of his family, their characteristics, antecedents, and education, etc.; the kind of soil on which the wheat was grown, its physical and chemical characteristics, the weather prevailing during the growing season; the personal characteristics of the farmers growing the wheat and of the consumers who will ultimately use it; and so on indefinitely. Any attempt to move very far in achieving this kind of “realism” is certain to render a theory utterly useless.

Of course, the notion of a completely realistic theory is in part a straw man. No critic of a theory would accept this logical extreme as his objective; he would say that the “assumptions” of the theory being criticized were

“too” unrealistic and that his objective was a set of assumptions that were “more” realistic though still not completely and slavishly so. But so long as the test of “realism” is the directly perceived descriptive accuracy of the “assumptions” – for example, the observation that “businessmen do not appear to be either as avaricious or as dynamic or as logical as marginal theory portrays them”²³ or that “it would be utterly impractical under present conditions for the manager of a multiprocess plant to attempt . . . to work out and equate marginal costs and marginal revenues for each productive factor”²⁴ – there is no basis for making such a distinction, that is, for stopping short of the straw man depicted in the preceding paragraph. What is the criterion by which to judge whether a particular departure from realism is or is not acceptable? Why is it more “unrealistic” in analyzing business behavior to neglect the magnitude of businessmen’s costs than the color of their eyes? The obvious answer is because the first makes more difference to business behavior than the second; but there is no way of knowing that this is so simply by observing that businessmen do have costs of different magnitudes and eyes of different color. Clearly it can only be known by comparing the effect on the discrepancy between actual and predicted behavior of taking the one factor or the other into account. Even the most extreme proponents of realistic assumptions are thus necessarily driven to reject their own criterion and to accept the test by prediction when they classify alternative assumptions as more or less realistic.²⁵

The basic confusion between descriptive accuracy and analytical relevance that underlies most criticisms of economic theory on the grounds that its assumptions are unrealistic as well as the plausibility of the views that lead to this confusion are both strikingly illustrated by a seemingly innocuous remark in an article on business-cycle theory that “economic phenomena are varied and complex, so any comprehensive theory of the business cycle that can apply closely to reality must be very complicated.”²⁶ A fundamental hypothesis of science is that appearances are deceptive and that there is a way of looking at or interpreting or organizing the evidence that will reveal superficially disconnected and diverse phenomena to be manifestations of a more fundamental and relatively simple structure. And the test of this hypothesis, as of any other, is its fruits – a test that science has so far met with dramatic success. If a class of “economic phenomena” appears varied and complex, it is, we must suppose, because we have no adequate theory to explain them. Known facts cannot be set on one side; a theory to apply “closely to reality,” on the other. A theory is the way we perceive “facts,” and we cannot perceive “facts” without a theory. Any assertion that economic phenomena *are* varied and complex denies the tentative state of knowledge

that alone makes scientific activity meaningful; it is in a class with John Stuart Mill's justly ridiculed statement that "happily, there is nothing in the laws of value which remains [1848] for the present or any future writer to clear up; the theory of the subject is complete."²⁷

The confusion between descriptive accuracy and analytical relevance has led not only to criticisms of economic theory on largely irrelevant grounds but also to misunderstanding of economic theory and misdirection of efforts to repair supposed defects. "Ideal types" in the abstract model developed by economic theorists have been regarded as strictly descriptive categories intended to correspond directly and fully to entities in the real world independently of the purpose for which the model is being used. The obvious discrepancies have led to necessarily unsuccessful attempts to construct theories on the basis of categories intended to be fully descriptive.

This tendency is perhaps most clearly illustrated by the interpretation given to the concepts of "perfect competition" and "monopoly" and the development of the theory of "monopolistic" or "imperfect competition." Marshall, it is said, assumed "perfect competition"; perhaps there once was such a thing. But clearly there is no longer, and we must therefore discard his theories. The reader will search long and hard – and I predict unsuccessfully – to find in Marshall any explicit assumption about perfect competition or any assertion that in a descriptive sense the world is composed of atomistic firms engaged in perfect competition. Rather, he will find Marshall saying: "At one extreme are world markets in which competition acts directly from all parts of the globe; and at the other those secluded markets in which all direct competition from afar is shut out, though indirect and transmitted competition may make itself felt even in these; and about midway between these extremes lie the great majority of the markets which the economist and the business man have to study."²⁸ Marshall took the world as it is; he sought to construct an "engine" to analyze it, not a photographic reproduction of it.

In analyzing the world as it is, Marshall constructed the hypothesis that, for many problems, firms could be grouped into "industries" such that the similarities among the firms in each group were more important than the differences among them. These are problems in which the important element is that a group of firms is affected alike by some stimulus – a common change in the demand for their products, say, or in the supply of factors. But this will not do for all problems: the important element for these may be the differential effect on particular firms.

The abstract model corresponding to this hypothesis contains two "ideal" types of firms: atomistically competitive firms, grouped into industries, and

monopolistic firms. A firm is competitive if the demand curve for its output is infinitely elastic with respect to its own price for some price and all outputs, given the prices charged by all other firms; it belongs to an “industry” defined as a group of firms producing a single “product.” A “product” is defined as a collection of units that are perfect substitutes to purchasers so the elasticity of demand for the output of one firm with respect to the price of another firm in the same industry is infinite for some price and some outputs. A firm is monopolistic if the demand curve for its output is not infinitely elastic at some price for all outputs.²⁹ If it is a monopolist, the firm is the industry.³⁰

As always, the hypothesis as a whole consists not only of this abstract model and its ideal types but also of a set of rules, mostly implicit and suggested by example, for identifying actual firms with one or the other ideal type and for classifying firms into industries. The ideal types are not intended to be descriptive; they are designed to isolate the features that are crucial for a particular problem. Even if we could estimate directly and accurately the demand curve for a firm’s product, we could not proceed immediately to classify the firm as perfectly competitive or monopolistic according as the elasticity of the demand curve is or is not infinite. No observed demand curve will ever be precisely horizontal, so the estimated elasticity will always be finite. The relevant question always is whether the elasticity is “sufficiently” large to be regarded as infinite, but this is a question that cannot be answered, once for all, simply in terms of the numerical value of the elasticity itself, any more than we can say, once for all, whether an air pressure of 15 pounds per square inch is “sufficiently” close to zero to use the formula $s = 1/2 \text{ } g t^2$. Similarly, we cannot compute cross-elasticities of demand and then classify firms into industries according as there is a “substantial gap in the cross-elasticities of demand.” As Marshall says, “The question where the lines of division between different commodities [i.e., industries] should be drawn must be settled by convenience of the particular discussion.”³¹ Everything depends on the problem; there is no inconsistency in regarding the same firm as if it were a perfect competitor for one problem, and a monopolist for another, just as there is none in regarding the same chalk mark as a Euclidean line for one problem, a Euclidean surface for a second, and a Euclidean solid for a third. The size of the elasticity and cross-elasticity of demand, the number of firms producing physically similar products, etc., are all relevant because they are or may be among the variables used to define the correspondence between the ideal and real entities in a particular problem and to specify the circumstances under which the theory holds sufficiently well; but they do not provide, once for all, a classification of firms as competitive or monopolistic.

An example may help to clarify this point. Suppose the problem is to determine the effect on retail prices of cigarettes of an increase, expected to be permanent, in the federal cigarette tax. I venture to predict that broadly correct results will be obtained by treating cigarette firms as if they were producing an identical product and were in perfect competition. Of course, in such a case, "some convention must be made as to the" number of Chesterfield cigarettes "which are taken as equivalent" to a Marlboro.³²

On the other hand, the hypothesis that cigarette firms would behave as if they were perfectly competitive would have been a false guide to their reactions to price control in World War II, and this would doubtless have been recognized before the event. Costs of the cigarette firms must have risen during the war. Under such circumstances perfect competitors would have reduced the quantity offered for sale at the previously existing price. But, at that price, the wartime rise in the income of the public presumably increased the quantity demanded. Under conditions of perfect competition strict adherence to the legal price would therefore imply not only a "shortage" in the sense that quantity demanded exceeded quantity supplied but also an absolute decline in the number of cigarettes produced. The facts contradict this particular implication: there was reasonably good adherence to maximum cigarette prices, yet the quantities produced increased substantially. The common force of increased costs presumably operated less strongly than the disruptive force of the desire by each firm to keep its share of the market, to maintain the value and prestige of its brand name, especially when the excess-profits tax shifted a large share of the costs of this kind of advertising to the government. For this problem the cigarette firms cannot be treated *as if* they were perfect competitors.

Wheat farming is frequently taken to exemplify perfect competition. Yet, while for some problems it is appropriate to treat cigarette producers as if they comprised a perfectly competitive industry, for some it is not appropriate to treat wheat producers as if they did. For example, it may not be if the problem is the differential in prices paid by local elevator operators for wheat.

Marshall's apparatus turned out to be most useful for problems in which a group of firms is affected by common stimuli, and in which the firms can be treated *as if* they were perfect competitors. This is the source of the misconception that Marshall "assumed" perfect competition in some descriptive sense. It would be highly desirable to have a more general theory than Marshall's, one that would cover at the same time both those cases in which differentiation of product or fewness of numbers makes an essential difference and those in which it does not. Such a theory would enable us

to handle problems we now cannot and, in addition, facilitate determination of the range of circumstances under which the simpler theory can be regarded as a good enough approximation. To perform this function, the more general theory must have content and substance; it must have implications susceptible to empirical contradiction and of substantive interest and importance.

The theory of imperfect or monopolistic competition developed by Chamberlin and Robinson is an attempt to construct such a more general theory.³³ Unfortunately, it possesses none of the attributes that would make it a truly useful general theory. Its contribution has been limited largely to improving the exposition of the economics of the individual firm and thereby the derivation of implications of the Marshallian model, refining Marshall's monopoly analysis, and enriching the vocabulary available for describing industrial experience.

The deficiencies of the theory are revealed most clearly in its treatment of, or inability to treat, problems involving groups of firms – Marshallian “industries.” So long as it is insisted that differentiation of product is essential – and it is the distinguishing feature of the theory that it does insist on this point – the definition of an industry in terms of firms producing an identical product cannot be used. By that definition each firm is a separate industry. Definition in terms of “close” substitutes or a “substantial” gap in cross-elasticities evades the issue, introduces fuzziness and undefinable terms into the abstract model where they have no place, and serves only to make the theory analytically meaningless – “close” and “substantial” are in the same category as a “small” air pressure.³⁴ In one connection Chamberlin implicitly defines an industry as a group of firms having identical cost and demand curves.³⁵ But this, too, is logically meaningless so long as differentiation of product is, as claimed, essential and not to be put aside. What does it mean to say that the cost and demand curves of a firm producing bulldozers are identical with those of a firm producing hairpins?³⁶ And if it is meaningless for bulldozers and hairpins, it is meaningless also for two brands of toothpaste – so long as it is insisted that the difference between the two brands is fundamentally important.

The theory of monopolistic competition offers no tools for the analysis of an industry and so no stopping place between the firm at one extreme and general equilibrium at the other.³⁷ It is therefore incompetent to contribute to the analysis of a host of important problems: the one extreme is too narrow to be of great interest; the other, too broad to permit meaningful generalizations.³⁸

VI. Conclusion

Economics as a positive science is a body of tentatively accepted generalizations about economic phenomena that can be used to predict the consequences of changes in circumstances. Progress in expanding this body of generalizations, strengthening our confidence in their validity, and improving the accuracy of the predictions they yield is hindered not only by the limitations of human ability that impede all search for knowledge but also by obstacles that are especially important for the social sciences in general and economics in particular, though by no means peculiar to them. Familiarity with the subject matter of economics breeds contempt for special knowledge about it. The importance of its subject matter to everyday life and to major issues of public policy impedes objectivity and promotes confusion between scientific analysis and normative judgment. The necessity of relying on uncontrolled experience rather than on controlled experiment makes it difficult to produce dramatic and clear-cut evidence to justify the acceptance of tentative hypotheses. Reliance on uncontrolled experience does not affect the fundamental methodological principle that a hypothesis can be tested only by the conformity of its implications or predictions with observable phenomena; but it does render the task of testing hypotheses more difficult and gives greater scope for confusion about the methodological principles involved. More than other scientists, social scientists need to be self-conscious about their methodology.

One confusion that has been particularly rife and has done much damage is confusion about the role of “assumptions” in economic analysis. A meaningful scientific hypothesis or theory typically asserts that certain forces are, and other forces are not, important in understanding a particular class of phenomena. It is frequently convenient to present such a hypothesis by stating that the phenomena it is desired to predict behave in the world of observation *as if* they occurred in a hypothetical and highly simplified world containing only the forces that the hypothesis asserts to be important. In general, there is more than one way to formulate such a description – more than one set of “assumptions” in terms of which the theory can be presented. The choice among such alternative assumptions is made on the grounds of the resulting economy, clarity, and precision in presenting the hypothesis; their capacity to bring indirect evidence to bear on the validity of the hypothesis by suggesting some of its implications that can be readily checked with observation or by bringing out its connection with other hypotheses dealing with related phenomena; and similar considerations.

Such a theory cannot be tested by comparing its “assumptions” directly with “reality.” Indeed, there is no meaningful way in which this can be done. Complete “realism” is clearly unattainable, and the question whether a theory is realistic “enough” can be settled only by seeing whether it yields predictions that are good enough for the purpose in hand or that are better than predictions from alternative theories. Yet the belief that a theory can be tested by the realism of its assumptions independently of the accuracy of its predictions is widespread and the source of much of the perennial criticism of economic theory as unrealistic. Such criticism is largely irrelevant, and, in consequence, most attempts to reform economic theory that it has stimulated have been unsuccessful.

The irrelevance of so much criticism of economic theory does not of course imply that existing economic theory deserves any high degree of confidence. These criticisms may miss the target, yet there may be a target for criticism. In a trivial sense, of course, there obviously is. Any theory is necessarily provisional and subject to change with the advance of knowledge. To go beyond this platitude, it is necessary to be more specific about the content of “existing economic theory” and to distinguish among its different branches; some parts of economic theory clearly deserve more confidence than others. A comprehensive evaluation of the present state of positive economics, summary of the evidence bearing on its validity, and assessment of the relative confidence that each part deserves is clearly a task for a treatise or a set of treatises, if it be possible at all, not for a brief paper on methodology.

About all that is possible here is the cursory expression of a personal view. Existing relative price theory, which is designed to explain the allocation of resources among alternative ends and the division of the product among the co-operating resources and which reached almost its present form in Marshall’s *Principles of Economics*, seems to me both extremely fruitful and deserving of much confidence for the kind of economic system that characterizes Western nations. Despite the appearance of considerable controversy, this is true equally of existing static monetary theory, which is designed to explain the structural or secular level of absolute prices, aggregate output, and other variables for the economy as a whole and which has had a form of the quantity theory of money as its basic core in all of its major variants from David Hume to the Cambridge School to Irving Fisher to John Maynard Keynes. The weakest and least satisfactory part of current economic theory seems to me to be in the field of monetary dynamics, which is concerned with the process of adaptation of the economy as a whole to changes in conditions and so with short-period fluctuations in aggregate activity. In this

field we do not even have a theory that can appropriately be called “the” existing theory of monetary dynamics.

Of course, even in relative price and static monetary theory there is enormous room for extending the scope and improving the accuracy of existing theory. In particular, undue emphasis on the descriptive realism of “assumptions” has contributed to neglect of the critical problem of determining the limits of validity of the various hypotheses that together constitute the existing economic theory in these areas. The abstract models corresponding to these hypotheses have been elaborated in considerable detail and greatly improved in rigor and precision. Descriptive material on the characteristics of our economic system and its operations have been amassed on an unprecedented scale. This is all to the good. But, if we are to use effectively these abstract models and this descriptive material, we must have a comparable exploration of the criteria for determining what abstract model it is best to use for particular kinds of problems, what entities in the abstract model are to be identified with what observable entities, and what features of the problem or of the circumstances have the greatest effect on the accuracy of the predictions yielded by a particular model or theory.

Progress in positive economics will require not only the testing and elaboration of existing hypotheses but also the construction of new hypotheses. On this problem there is little to say on a formal level. The construction of hypotheses is a creative act of inspiration, intuition, invention; its essence is the vision of something new in familiar material. The process must be discussed in psychological, not logical, categories; studies in autobiographies and biographies, not treatises on scientific method; and promoted by maxim and example, not syllogism or theorem.

Notes

1. (London: Macmillan & Co., 1981), pp. 34–5 and 46.
2. Social science or economics is by no means peculiar in this respect – witness the importance of personal beliefs and of “home” remedies in medicine wherever obviously convincing evidence for “expert” opinion is lacking. The current prestige and acceptance of the views of physical scientists in their fields of specialization – and, all too often, in other fields as well – derives, not from faith alone, but from the evidence of their works, the success of their predictions, and the dramatic achievements from applying their results. When economics seemed to provide such evidence of its worth, in Great Britain in the first half of the nineteenth century, the prestige and acceptance of “scientific economics” rivaled the current prestige of the physical sciences.
3. The interaction between the observer and the process observed that is so prominent a feature of the social sciences, besides its more obvious parallel in the

physical sciences, has a more subtle counterpart in the indeterminacy principle arising out of the interaction between the process of measurement and the phenomena being measured. And both have a counterpart in pure logic in Gödel's theorem, asserting the impossibility of a comprehensive self-contained logic. It is an open question whether all three can be regarded as different formulations of an even more general principle.

4. One rather more complex example is stabilization policy. Superficially, divergent views on this question seem to reflect differences in objectives; but I believe that this impression is misleading and that at bottom the different views reflect primarily different judgments about the source of fluctuations in economic activity and the effect of alternative countercyclical action. For one major positive consideration that accounts for much of the divergence see "The Effects of a Full-Employment Policy on Economic Stability: A Formal Analysis," *infra*, pp. 117–32. For a summary of the present state of professional views on this question see "The Problem of Economic Instability," a report of a subcommittee of the Committee on Public Issues of the American Economic Association, *American Economic Review*, XL (September, 1950), 501–38.
5. Final quoted phrase from Alfred Marshall, "The Present Position of Economics" (1885), reprinted in *Memorials of Alfred Marshall*, ed. A. C. Pigou (London: Macmillan & Co., 1925), p. 164. See also "The Marshallian Demand Curve," *infra*, pp. 56–7, 90–1.
6. See "Lange on Price Flexibility and Employment: A Methodological Criticism," *infra*, pp. 282–9.
7. "The Marshallian Demand Curve," *infra*, p. 57.
8. The qualification is necessary because the "evidence" may be internally contradictory, so there may be no hypothesis consistent with it. See also "Lange on Price Flexibility and Employment," *infra*, pp. 282–3.
9. See "Lange on Price Flexibility and Employment," *infra*, *passim*.
10. See also Milton Friedman and L. J. Savage, "The Expected-Utility Hypothesis and the Measurability of Utility," *Journal of Political Economy*, LX (December, 1952), 463–74, esp. pp. 465–7.
11. In recent years some economists, particularly a group connected with the Cowles Commission for Research in Economics at the University of Chicago, have placed great emphasis on a division of this step of selecting a hypothesis consistent with known evidence into two substeps: first, the selection of a class of admissible hypotheses from all possible hypotheses (the choice of a "model" in their terminology); second, the selection of one hypothesis from this class (the choice of a "structure"). This subdivision may be heuristically valuable in some kinds of work, particularly in promoting a systematic use of available statistical evidence and theory. From a methodological point of view, however, it is an entirely arbitrary subdivision of the process of deciding on a particular hypothesis that is on a par with many other subdivisions that may be convenient for one purpose or another or that may suit the psychological needs of particular investigators.

One consequence of this particular subdivision has been to give rise to the so-called "identification" problem. As noted above, if one hypothesis is consistent with available evidence, an infinite number are. But while this is true for the class of hypotheses as a whole, it may not be true of the subclass obtained

in the first of the above two steps – the “model.” It may be that the evidence to be used to select the final hypothesis from the subclass can be consistent with at most one hypothesis in it, in which case the “model” is said to be “identified”; otherwise it is said to be “unidentified.” As is clear from this way of describing the concept of “identification,” it is essentially a special case of the more general problem of selecting among the alternative hypotheses equally consistent with the evidence – a problem that must be decided by some such arbitrary principle as Occam’s razor. The introduction of two substeps in selecting a hypothesis makes this problem arise at the two corresponding stages and gives it a special cast. While the class of all hypotheses is always unidentified, the subclass in a “model” need not be, so the problem arises of conditions that a “model” must satisfy to be identified. However useful the two substeps may be in some contexts, their introduction raises the danger that different criteria will unwittingly be used in making the same kind of choice among alternative hypotheses at two different stages.

On the general methodological approach discussed in this footnote see Tryvge Haavelmo, “The Probability Approach in Econometrics,” *Econometrica*, Vol. XII (1944), Supplement; Jacob Marschak, “Economic Structure, Path, Policy, and Prediction,” *American Economic Review*, XXXVII, (May, 1947), 81–84, and “Statistical Inference in Economics: An Introduction,” in T. C. Koopmans (ed.), *Statistical Inference in Dynamic Economic Models* (New York: John Wiley & Sons, 1950); T. C. Koopmans, “Statistical Estimation of Simultaneous Economic Relations,” *Journal of the American Statistical Association*, XL (December, 1945), 448–66; Gershon Cooper, “The Role of Economic Theory in Econometric Models,” *Journal of Farm Economics*, XXX (February, 1948), 101–16. On the identification problem see Koopmans, “Identification Problems in Econometric Model Construction,” *Econometrica*, XVII (April, 1949), 125–44; Leonid Hurwicz, “Generalization of the Concept of Identification,” in Koopmans (ed.), *Statistical Inference in Dynamic Economic Models*.

12. The converse of the proposition does not of course hold: assumptions that are unrealistic (in this sense) do not guarantee a significant theory.
13. See R. A. Lester, “Shortcomings of Marginal Analysis for Wage-Employment Problems,” *American Economic Review*, XXXVI (March, 1946), 62–82; Fritz Machlup, “Marginal Analysis and Empirical Research,” *American Economic Review*, XXXVI (September, 1946), 519–54; R. A. Lester, “Marginalism, Minimum Wages, and Labor Markets,” *American Economic Review*, XXXVII (March, 1947), 135–48; Fritz Machlup, “Rejoinder to an Antimarginalist,” *American Economic Review*, XXXVII (March, 1947), 148–54; G. J. Stigler, “Professor Lester and the Marginalists,” *American Economic Review*, XXXVII (March, 1947), 154–57; H. M. Oliver, Jr., “Marginal Theory and Business Behavior,” *American Economic Review*, XXXVII (June, 1947), 375–83; R. A. Gordon, “Short-Period Price Determination in Theory and Practice,” *American Economic Review*, XXXVIII (June, 1948), 265–88.

It should be noted that, along with much material purportedly bearing on the validity of the “assumptions” of marginal theory, Lester does refer to evidence on the conformity of experience with the implications of the theory, citing the reactions of employment in Germany to the Papen plan and in the United States

to changes in minimum-wage legislation as examples of lack of conformity. However, Stigler's brief comment is the only one of the other papers that refers to this evidence. It should be noted that Machlup's thorough and careful exposition of the logical structure and meaning of marginal analysis is called for by the misunderstandings on this score that mar Lester's paper and almost conceal the evidence he presents that is relevant to the key issue he raises. But, in Machlup's emphasis on the logical structure, he comes perilously close to presenting the theory as a pure tautology, though it is evident at a number of points that he is aware of this danger and anxious to avoid it. The papers by Oliver and Gordon are the most extreme in the exclusive concentration on the conformity of the behavior of businessmen with the "assumptions" of the theory.

14. This example, and some of the subsequent discussion, though independent in origin, is similar to and in much the same spirit as an example and the approach in an important paper by Armen A. Alchian, "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy*, LVIII (June, 1950), 211–21.
15. Milton Friedman and L. J. Savage, "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*, LVI (August, 1948), 298. Reprinted in American Economic Association, *Readings in Price Theory* (Chicago: Richard D. Irwin, Inc., 1952), pp. 57–96.
16. It seems better to use the term "profits" to refer to the difference between actual and "expected" results, between *ex post* and *ex ante* receipts. "Profits" are then a result of uncertainty and, as Alchian (*op. cit.*, p. 212), following Tintner, points out, cannot be deliberately maximized in advance. Given uncertainty, individuals or firms choose among alternative anticipated probability distributions of receipts or incomes. The specific content of a theory of choice among such distributions depends on the criteria by which they are supposed to be ranked. One hypothesis supposes them to be ranked by the mathematical expectation of utility corresponding to them (see Friedman and Savage, "The Expected-Utility Hypothesis and the Measurability of Utility," *op. cit.*). A special case of this hypothesis or an alternative to it ranks probability distribution by the mathematical expectation of the money receipts corresponding to them. The latter is perhaps more applicable, and more frequently applied, to firms than to individuals. The term "expected returns" is intended to be sufficiently broad to apply to any of these alternatives.

The issues alluded to in this note are not basic to the methodological issues being discussed, and so are largely by-passed in the discussion that follows.

17. See George J. Stigler, "A Theory of Delivered Price Systems," *American Economic Review*, XXXIX (December, 1949), 1143–57.
18. See Friedman and Savage, "The Expected-Utility Hypothesis and the Measurability of Utility," *op. cit.*, pp. 466–7, for another specific example of this kind of indirect test.
19. A rigorous statement of this hypothesis would of course have to specify how "extent of racial or religious discrimination" and "degree of monopoly" are to be judged. The loose statement in the text is sufficient, however, for present purposes.
20. Thorstein Veblen, "Why Is Economics Not an Evolutionary Science?" (1898), reprinted in *The Place of Science in Modern Civilization* (New York, 1919), p. 73.

21. Oliver, *op. cit.*, p. 381.
22. See H. D. Henderson, "The Significance of the Rate of Interest," *Oxford Economic Papers*, No. 1 (October, 1938), pp. 1–13; J. E. Meade and P. W. S. Andrews, "Summary of Replies to Questions on Effects of Interest Rates," *Oxford Economic Papers*, No. 1 (October, 1938), pp. 14–31; R. F. Harrod, "Price and Cost in Entrepreneurs' Policy," *Oxford Economic Papers*, No. 2 (May, 1939), pp. 1–11; and R. J. Hall and C. J. Hitch, "Price Theory and Business Behavior," *Oxford Economic Papers*, No. 2 (May, 1939), pp. 12–45; Lester, "Shortcomings of Marginal Analysis for Wage-Employment Problems," *op. cit.*; Gordon, *op. cit.* See Fritz Machlup, "Marginal Analysis and Empirical Research," *op. cit.*, esp. Sec. II, for detailed criticisms of questionnaire methods.
I do not mean to imply that questionnaire studies of businessmen's or others' motives or beliefs about the forces affecting their behavior are useless for all purposes in economics. They may be extremely valuable in suggesting leads to follow in accounting for divergencies between predicted and observed results; that is, in constructing new hypotheses or revising old ones. Whatever their suggestive value in this respect, they seem to me almost entirely useless as a means of testing the validity of economic hypotheses. See my comment on Albert G. Hart's paper, "Liquidity and Uncertainty," *American Economic Review*, XXXIX (May, 1949), 198–99.
23. Oliver, *op. cit.*, p. 382.
24. Lester, "Shortcomings of Marginal Analysis for Wage-Employment Problems," *op. cit.*, p. 75.
25. E.g., Gordon's direct examination of the "assumptions" leads him to formulate the alternative hypothesis generally favored by the critics of the maximization-of-returns hypotheses follows: "There is an irresistible tendency to price on the basis of average total costs for some 'normal' level of output. This is the yardstick, the short-cut, that businessmen and accountants use, and their aim is more to earn satisfactory profits and play safe than to maximize profits" (*op. cit.*, p. 275). Yet he essentially abandons this hypothesis, or converts it into a tautology, and in the process implicitly accepts the test by prediction when he later remarks: "Full cost and satisfactory profits may continue to be the objectives even when total costs are shaded to meet competition or exceeded to take advantage of a sellers' market" (*ibid.*, p. 284). Where here is the "irresistible tendency"? What kind of evidence could contradict this assertion?
26. Sidney S. Alexander, "Issues of Business Cycle Theory Raised by Mr. Hicks," *American Economic Review*, XLI (December, 1951), 872.
27. *Principles of Political Economy* (Ashley ed.; Longmans, Green & Co., 1929), p. 436.
28. *Principles*, p. 329; see also pp. 35, 100, 341, 347, 375, 546.
29. This ideal type can be divided into two types: the oligopolistic firm, if the demand curve for its output is infinitely elastic at some price for some but not all outputs; the monopolistic firm proper, if the demand curve is nowhere infinitely elastic (except possibly at an output of zero).
30. For the oligopolist of the preceding note an industry can be defined as a group of firms producing the same product.
31. *Principles*, p. 100.
32. Quoted parts from *ibid.*

33. E. H. Chamberlin, *The Theory of Monopolistic Competition* (6th ed.; Cambridge: Harvard University Press, 1950); Joan Robinson, *The Economics of Imperfect Competition* (London: Macmillan & Co., 1933).
34. See R. L. Bishop, "Elasticities, Cross-elasticities, and Market Relationships," *American Economic Review*, XLII (December, 1952), 779–803, for a recent attempt to construct a rigorous classification of market relationships along these lines. Despite its ingenuity and sophistication, the result seems to me thoroughly unsatisfactory. It rests basically on certain numbers being classified as "large" or "small," yet there is no discussion at all of how to decide whether a particular number is "large" or "small," as of course there cannot be on a purely abstract level.
35. *Op. cit.*, p. 82.
36. There always exists a transformation of quantities that will make either the cost curves or the demand curves identical; this transformation need not, however, be linear, in which case it will involve different-sized units of one product at different levels of output. There does not necessarily exist a transformation that will make both pairs of curves identical.
37. See Robert Triffin, *Monopolistic Competition and General Equilibrium Theory* (Cambridge: Harvard University Press, 1940), esp. pp. 188–89.
38. For a detailed critique see George J. Stigler, "Monopolistic Competition in Retrospect," in *Five Lectures on Economic Problems* (London: Macmillan & Co., 1949), pp. 12–14.