

CONSTRUCTING SOCIAL THEORIES

CHAPTER TWO

The Logic of Scientific Inference

I / FUNDAMENTAL FORMS OF SCIENTIFIC INFERENCE

In order to construct theories for a science, we must have in mind the logical requirements for testing the theories against the facts. Hence our first task is to outline the fundamental logical forms of scientific inference, or of "induction," which form the common basis of the sciences.

Theoretical and Empirical Statements

Scientific inference starts with a *theoretical statement*, an element of a theory, which says that one class of phenomena will be connected in a certain way with another class of phenomena. A famous example in sociological theory would be Durkheim's theory of egoistic suicide, which might be stated as: "A higher degree of individualism in a social group causes a higher rate of suicide in that group." Here "individualism" is a variable meaning, roughly, "the degree to which all activities of the person are controlled by

well-defined norms enforced regularly and effectively by people in the environment, as opposed to morality determined by the individual himself by his own decisions." A person is in a more individualistic social situation, then, when fewer demands are made on him—as when he is single rather than a married man with children, or when his country is not in crisis (so he has fewer duties). He is in a more individualistic situation also when a group of which he is a member does not govern his activity in detail but leaves it to his own discretion (as when he is a Protestant rather than a Catholic), when the group which regulates his conduct is not so compact that it can surround him with others who enforce the group's prescriptions (as, in Durkheim's time in France, the Jews were more compact than other social groups).

As we can see from this example, the concepts in the theoretical statement may be at quite different levels of abstraction, with "individualism" being a property of groups which is inferred quite indirectly while the suicide rate is at a low level of abstraction and is directly observable.

From this theoretical statement we derive, by logical deduction and by operational definitions of the concepts, an *empirical statement*. The theoretical statement then *implies* logically the empirical statement. An empirical statement is one which states that: "If we make such and such observations, they will have such and such results." For instance, some of the empirical statements which Durkheim derived from his theory of egoistic suicide are: Protestants in France will have higher suicide rates than Catholics in France; Protestant regions of German provinces will have higher rates of suicide than Catholic regions; married men in France will have lower rates of suicide than single men and will have even lower rates if they have children; men who practice the free professions and generally well-educated men will have higher rates of suicide than workers or less educated people; in times of parliamentary crisis, the suicide rates will go down in France; and others. Such logical derivation from theoretical statements involves stating the meaning of the concepts in terms of observations. The statements above, that Protestants have a higher degree of individualism than Catholics, bachelors than married men with children, and populations during parliamentary crises than populations during routine politics, are essential

in the derivation of empirical statements. And derivation involves straightforward logical deduction.

After this logical deduction of empirical statements from the theory, one can make the observations called for in the empirical statements to see whether or not they are true. As a practical matter, it is important to describe the observations in such a way that they can actually be made with the resources the investigator has at his disposal, but that does not affect the logic of the matter. All theories will imply some empirical statements (that is, some descriptions of possible observations) which cannot actually be tested because of lack of time, or lack of money, or technical impossibility.

Testing Theories with Observations

Once the observations are in fact made, we can compare them with the empirical statement and find out whether or not the statement is true. For simplicity in the discussion, let us call the theory A , and one of its empirical consequences B . And let us use an arrow with a double shaft \Rightarrow for "implies." Then we have two logical situations:

SITUATION I	SITUATION II
$A \Rightarrow B$	$A \Rightarrow B$
B false	B true
A false	A more credible

In situation I, classical logic gives the results. If A implies B , then not- B implies not- A . If Durkheim's theory implied that Protestants ought to have a higher rate of suicide in France, and they in fact have a lower rate, then his theory is false. (The difficulty may be either in the statement that individualism causes suicide or in the statement that Protestants are more individualistic than Catholics, but one or the other must be wrong.) The deduction could be wrong for "irrelevant" reasons, as for instance because there are not enough observations to give a good estimate of the "true" suicide rate and therefore the theory should have been stated explicitly with a restriction on the sizes of the groups to which

The notation here follows Georgi W. Polya, *Patterns of Plausible Inference*, Vol. 2 of *Mathematics and Plausible Reasoning* (Princeton, N.J.: Princeton University Press, 1954).

it applies (but such complexities will be dealt with later). The canons of logic demand that we reject our theory if it implies something that is false.

But classical canons of logic have nothing to say in the second case, if our observations "confirm" or "support" or "are consistent with" our theory. Yet our intuition tells us strongly that *something* has happened to our theory. We have, as we say, "tested" the theory against the facts, and it has stood the test. Intuitively it seems to us that by virtue of the test our theory has become more believable or, as we have said above, "A more credible." The whole edifice of scientific inference rests on the logical situation described above in situation II, on "affirming the consequent" of a theory, in the language of logic.

Intuitively we can formulate the logical incompleteness of situation II, the situation of affirming the consequent, by saying that "there are a lot of other possible explanations for *B*." For instance, the higher suicide rate of French Protestants might be explained by their occupations, by the lesser emphasis on the sin of suicide in Protestant theology, by the fact that confessors are available to every Catholic in times of trouble and distress, and so forth. Yet we still feel that if Durkheim's theory implied this fact, and if the fact turned out the way it was predicted, Durkheim's theory has been shown to be a better theory than it was before.²

Multiple Tests of Theories

Our problem is to analyze this intuition on which science is based in order to describe more precisely how facts can "support theories" and hence what kinds of facts and what kinds of theories we need to be oriented to. First, the list of implications given above

²Sometimes the dependence of scientific inference on situation II is formulated by saying that "the test of scientific theories is prediction." A theory may, however, be able to predict things and yet be false, as shown by the falsification of *other* "predictions." Moreover, various systematic ways of guessing the future values of some variable ("projections") are also sometimes called predictions, as in the statement, "The Protestant suicide rate in France will probably be higher than the Catholic rate next year, since it was higher last year, and experience shows a high degree of stability of suicide rates." Such projections are often administratively useful but have little to do with science.

from Durkheim's theory of egoistic suicide suggests an elaboration. For the theory implies that Protestants will have higher suicide rates than Catholics, and hence that (B_1) Protestant countries will have higher suicide rates than Catholic countries, that (B_2) Protestant regions of Germany will have higher suicide rates than Catholic regions, and that (B_3) Protestants in France will have higher suicide rates than Catholics. Durkheim makes these derivations and collects the statistics to test his theory. The logical alternatives then become:

SITUATION I	SITUATION II	SITUATION III
$A \Rightarrow B$	$A \Rightarrow B$	$A \Rightarrow B_1, B_2, B_3$
<u>B false</u>	<u>B true</u>	<u>B_1, B_2, B_3 all true</u>
A false	A more credible	A substantially more credible

That is, a multiple test of a theory is more convincing than a single test. Intuitively speaking, we have given the theory more chances to be disproved, and it has stood up under them all. Presumably there are fewer other possible explanations of B_1 , B_2 , and B_3 taken together than there are for B as an isolated empirical statement.

Now let us suppose that we have already derived B_1 , B_2 , and B_3 from Durkheim's theory and have found that indeed Protestant countries, Protestant regions in Germany, and Protestants in France have lower suicide rates. Then let us consider the credibility added to Durkheim's theory if we were to show that the various regions of Austria had higher suicide rates, the higher their proportion Protestant. Compare this with the credibility added if we show that men with children have a lower suicide rate than bachelors and men without children. That is, suppose that we can make only one more observation, given our resources, and we have to choose between an empirical statement very similar to those we have already proved and one very different. It is clear that the "surprise value" of the observation on bachelors is much greater than the surprise value of *another* observation on Protestants and Catholics. If the theory can imply such different kinds of empirical statements as one about religion and one about marital status, then we feel intuitively that the theory has been subjected to a tougher test than it has if we merely repeat observations similar to those we have already made.

SITUATION I	SITUATION II	SITUATION III	SITUATION IV
$A \Rightarrow B$	$A \Rightarrow B$	$A \Rightarrow B_1, B_2, B_3$	$A \Rightarrow B_1, B_2, B_3$
<u>B false</u>	<u>B true</u>	<u>B_1, B_2, B_3 similar</u>	<u>B_1, B_2, B_3 different</u>
A false	A more credible	A substantially more credible	A much more credible

This result is summarized in the figure above. It says both that the more different things we can derive (situation III), and the more different kinds of implications we can derive (situation IV), the stronger will be our test of the theory. If the theory stands up under a tougher test, it becomes more credible than it is if it stands up when we have subjected it only to weak tests. If it fails any of the tests, it is false, either in the underlying statement or in the specification of the observations which the concepts of the theory refer to.

The Fundamental Criterion of a Strong Test of Theory

What guides our intuition to the conclusion that multiple, different tests of the consequences of a theory are better than a single, isolated test? In order to answer this question, we must consider the *alternative theories* which might be explanations of various phenomena in the world. In general, we imagine that before our investigation there are many alternative theories that one might hold about what goes on in the world. Many of these will be theories that someone else has already thought of, and many of them will be theories that we will think of during our investigation. But many of the theories that might be true will be theories that no one has ever thought of. In other words, we have a very large class of possible theories which are consistent with past knowledge, some of which are known and some of which are not.

For any given observation which is an implication of A , say B_1 , there will be *some* of the possible alternative theories which will imply not- B_1 . If we then demonstrate B_1 , these alternative theories are falsified. This leaves us with *fewer alternative possible theories* to our own. We can diagram this situation as follows (this is exactly comparable to situation II above, but formulated in a different way):

SITUATION II, REFORMULATED

(Situation before testing B₁)

A or (C, D, E, . . . , Q, R, S, . . .)

 $A \Rightarrow B_1$ $(C, E, . . . , Q, S, . . .) \Rightarrow B_1$ $D, R \Rightarrow \text{not-}B_1$ B_1 true

 D, R false (by classical logic)

A or (C, E, . . . , Q, S, . . .)

A more credible

(but also C, E, . . . , Q, S, . . . more credible)

That is, with the test B_1 we eliminate D and R from among the possible alternatives, making *all* of those which imply B_1 more credible. Among those theories made more credible by this elimination of competitors is our own theory, A .

Now if we have several implications from our theory, but all are quite similar, we eliminate some particular and erratic alternatives to our theory. If, for instance, we test only the implication from Durkheim's theory that Protestants in France will kill themselves more often than Catholics in France, some peculiarity of the situation of Protestants in France or of their history (such as the fact that they are a very small minority, or that there was selective migration at the time of the exodus of the French Huguenots) might explain the facts, since all these particular alternative theories imply the empirical statement. If we check the other derivations about Protestant countries and about regions in Germany, these peculiar explanations or alternative theories are also eliminated. But as the number of *similar* tests to the theory increases, *the number of alternative theories each new test eliminates* becomes much smaller.

Then if we turn to a quite different implication of our original theory, for instance, that men with children should kill themselves less often (in France) than bachelors, it will eliminate quite a different set of alternative theories. Thus what we mean when we say that two implications of a theory are "quite different" is that *there is almost no overlap* between the theories that imply the one empirical statement and the theories that imply the other. For instance, the number of alternative theories which imply *both* that Protes-

tants will have more suicide in the different regions of Germany *and* that men with children will have more than bachelors is much smaller than the number which would imply *both* that Protestant regions in Germany will have higher suicide rates *and* that Protestants in France will have higher suicide rates. This is so because all theories which involve Protestantism as a cause will imply both of the second pair, but (probably) only one of the first pair. We can represent this situation in the following figure:

A OR (C, D, E, F, . . . , Q, R, S, T, . . .)
(Situation before either test)

SITUATION III, REFORMULATED	SITUATION IV, REFORMULATED
$A \Rightarrow B_1, B_2$	$A \Rightarrow B_1, B_2$
B_1, B_2 similar	B_1, B_2 very different
$C, R \Rightarrow$ not B_1	$C, R \Rightarrow$ not B_1
$C, S \Rightarrow$ not B_2	$D, E \Rightarrow$ not B_2
<u>B_1, B_2 true</u>	<u>B_1, B_2 true</u>
C, R, S false	C, R, D, E false
A or (D, E, F, . . . , Q, T, . . .)	A or (F, . . . , Q, S, T, . . .)
A substantially more credible	A much more credible

Thus, because the empirical statements B_1 and B_2 are both implied by C , either one of them is sufficient to eliminate it, and (at least in this respect) the second test, which rejected C again, was "wasted." Only S , a particular explanation for the result B_1 which could not explain B_2 , was eliminated by the second test, by B_2 . But when B_1 and B_2 are different, there is no (or very little) overlap in the theories which might explain them. Hence more alternative theories are eliminated by checking two very different consequences of our theory. The theory S , the particular explanation for the results of the first test, is not eliminated by the recourse to a completely different implication of the theory. But it would probably have less weight in our thinking because of the increased credibility of our original theory.

Thus the basic logical process of science is the elimination of alternative theories (both those we know and those we do not) by investigating as many of the empirical consequences of each theory as is practical, always trying for the greatest possible variety in the implications tested.

Two Very General Alternative Theories: Statistical Inference

We are now in a position to discuss the logical role that statistical inference plays in science. There are two alternative theories that we always have to regard as possible explanations of any given set of observations, both of which give rise to *random distributions of the observations*. The first is that the observations were produced by *the way we designed our study*, especially by the *sample of observations we chose to make* out of those we could have made. In statistical theory, the observations we could have made are called either the population or the universe. The second alternative theory that we always have to take account of is that the observations were produced by *a large set of small influences operating in different di-rections*.³

The implication of this is that we always want to reject evidence if it can be explained either by the design of the research or by a large number of small, unorganized causes. Many observations made in daily life constitute such a small sample that any results might be explained by the selection of the sample from out of the total population of observations. For instance, Durkheim might have observed a small town in France to see whether Protestants killed themselves more often than Catholics. Whether we regard this as a sample of France's population (which it is, though an inefficient one) or the total population to which the theory applies, it is still true that any observations we make will have so few cases

³ The original theory of random distribution was worked out for the second situation, of a large set of small causes without any internal organization or pattern, in the theory of errors of measurement. Gauss's notion was that errors of measurement (after the measuring instruments were perfected and checked) were the result of a large number of small forces, such as friction, perceptual errors, recording errors, and the like, which operated in different directions in the different measurements of a phenomenon. Recently, statistical theory has been mainly developed as a theory of samples from a population, so that some people are unable to understand how statistical inference applies when they have measured the whole population of observations to which their theory applies. This failure of the imagination very often happens among people whose orientation to statistics is for administrative, rather than scientific, purposes, in which they only want to know a parameter in an administratively defined population. The orientation to statistics taken here is roughly that called "Bayesian."

of suicide in them that they might give different results than a complete count of suicides in France. On the other hand, such a small number of suicides might have been produced by a large number of small causes of suicide. That is, a random distribution could explain our observations.

Since the theory of sampling and the theory of large numbers of small causes apply to so many phenomena, if they imply our observations and if our theory also implies them, then our theory is very little more believable than it was before.⁴

The theory of sampling and the theory of large numbers of small, unorganized causes are quite highly developed mathematically. By knowing how we collected the observations, or by knowing the number of observations, we can often say quite precisely what kinds of observations would be consistent with these two very general theories. If the observations we actually make are inconsistent with these derivations (in other words, if statistical theory implies not- B_1 , and we observe B_1), then we eliminate *these* alternative theories.

The branch of statistics which is called the "design of experiments" deals with the problem of deriving from our theory consequences *which will not be implied* by either of the statistical theories (which are mathematically the same). It tries to specify which observations, if they were made, would be implied by our theory but not by statistical theory.⁵

The Crucial Experiment

The logic which we applied to statistical inference can be generalized to all cases in which we have *explicitly formulated alternative theories*. In many cases we can specify some of the most important alternatives to our own theory. If we can specify the

Many people would say that it was no more believable than before, but often observations that are consistent with chance are inconsistent with those alternative theories that imply a strong tendency for the observations to be different from the random distribution. In the language of statistics, our observations establish a confidence interval which includes the null hypothesis. Theories which imply parameters outside the confidence interval are thereby "rejected," or at least rendered unlikely.

Another part of the design of experiments in statistics deals with *turning other possible theories into statistical theories*, by a technique called *randomization*. Turning other theories into a statistical theory and then disproving the statistical theory is, of course, not the only way of disproving them, and it has nothing to do with the refutation of the sampling and large-group-of-small-causes explanations.

most important of alternative theories, then it is very inefficient to test our theory by picking empirical consequences at random, with the hope that some of them will be inconsistent with the main alternative theories. The rational thing to do is to look for those consequences of our theory whose *negation* is implied by the alternatives. We look for consequences, B_j , of our theory which will give us the logical situation shown below:

SITUATION V: THE CRUCIAL EXPERIMENT

$$\begin{array}{l}
 A \text{ or } C \text{ or } (D, E, \dots) \\
 (D, E, \dots) \text{ unlikely} \\
 A \Rightarrow B, \\
 C \Rightarrow \text{not-}B_1 \\
 \hline
 B, \text{ true} \\
 \hline
 C \text{ false} \\
 A \text{ or } (D, E, \dots) [(D, E, \dots) \text{ unlikely}] \\
 A \text{ very much more credible}
 \end{array}$$

By eliminating the most likely alternative theory, we increase the credibility of our theory much more than we do by eliminating alternatives at random by checking consequences of our theory without thinking.

Such a consequence of a theory (or, to speak more precisely, of both theories) is called a *crucial experiment* and is a description of a set of observations which will decide between two alternative theories, both of which according to present knowledge are quite likely. The purpose of the design of experiments in statistical theory is to construct crucial experiments between any given theory and statistical theory. In this case, usually, sampling theory (or some other theory reduced to sampling theory by randomization) plays the logical role of C above, and the design-of-experiments statistical worker searches for a B_j by working out the mathematical consequences of C as compared with A . Thus what is called the design of experiments in statistics is a particular example (and a particularly difficult example to understand) of the finding of crucial experiments.

Examples of Crucial Experiments

For more substantive examples, we can again turn to Durkheim's *Suicide*. The most popular alternative theory at the time Durkheim

lived was that suicide was the result of mental illness (not further specified) or was caused by the same causes that caused mental illness. Durkheim reasoned that if this were the case, then the same populations that had high rates of mental illness ought to have high rates of suicide. Since, as it happened, the social causes of suicide that Durkheim was interested in did not, apparently, cause variations in rates of mental illness, *his* theory implied that the correlation between rates of mental illness and rates of suicide would be insignificant. Thus he could describe a set of observations (the relations between rates of mental illness and rates of suicide, for various regions) which would show one result (positive correlation) if mental illness caused suicide, and a different result (insignificant correlation) if social causes were operating. He then made these observations, and the correlation between mental illness rates and suicide rates was insignificant. This disproved the alternative theory (as it was stated) and made his theory much more credible.

Of course, only the theory that was explicitly posed as an alternative is disproved, and sometimes that theory can then be reformulated so as to be consistent with the new observations. In this case, for instance, it seems sensible to suppose that particular kinds of mental illness, such as depressive disorders, would be related to suicide. There is some evidence that depressive disorders, which account for a relatively small part of mental disease and hence do not affect the overall rates very much, are related to quite different social variables than is schizophrenia. Schizophrenia accounts for the bulk of the variation in overall mental disease rates. Thus it could very well be that a *different alternative theory* relating suicide and mental disease might be true, one relating suicide specifically to depressive mental disorders. Durkheim's data fundamentally show that suicide rates are not related to rates of schizophrenia.

Perhaps an even more elegant example from Durkheim is his analysis of the suicide rates of Jews. By this point in the monograph Durkheim has shown that urbanism, education, and employment in commerce are all positively related to suicide, which he has interpreted as evidence for his theory of individualism. Now, he argues, if there were some other explanation for these connections of urbanism, education, and commercial employment to suicide, they *should operate just as well in a highly solidary, but educated, urban, commercial group*. If, on the other hand, urbanism, education, and

commercial employment are only indicators of individualism as opposed to solidarity, then when we find them in a highly solidary group they will not cause suicide. Thus if Durkheim can find a highly solidary, but educated, commercial, urban group, he can put a decisive test, a crucial experiment, to choose between his theory of individualism and very many of the alternative theories which would still be compatible with the data.

He argued that such a group were the Jews of France at his time. Because of their minority status, they formed a compact group with detailed, highly ritualized norms of daily conduct that were thoroughly enforced on each individual because the individual was almost always among his coreligionists. But the Jews were thoroughly urban, commercial, and much more highly educated than other Frenchmen. Consequently, Durkheim's theory predicted a very low suicide rate. Almost all other theories which could explain the relation of urbanism, education, and commercial employment to suicide would predict a very high rate. Consequently, when he showed that Jews at his time in France had very low suicide rates, he made his theory of individualism very much more credible.

Thus in the statistical design of experiments, in Durkheim's examination of mental disorders as causes of suicide, and in Durkheim's study of the suicide rates of Jews in France, we have the same logical situation. In each case we have explicitly developed the competing theory at least to such a level that we can derive some consequences from it. In the statistical design of experiments, this alternative theory is a theory of random distributions. In the case of mental infirmities and suicide the competing theory is that suicide is caused by, or has the same causes as, mental disorders in general. In the case of the Jews' suicide rates, the competition is that urbanism, education, and commercial employment will have the same effects on suicide, whatever their relation to individualism and social solidarity.

Then we choose one of the consequences of the theory we are testing that would contradict the alternative theory, and we make observations appropriate to checking *that* consequence of our theory in preference to checking all the consequences as we happen to think of them. If we are clever enough in inventing the alternative theories to make them inherently likely, and in deriving the consequences which will decide between them, we increase the effi-

ciency of our observation for the advance of science very greatly. Because the advance of science is much more economical when we can explicitly eliminate the most likely alternative theories, and because formulating the alternative theories and deriving their consequences is preeminently a theoretical task, the central gift of the great methodologist is his facility at formulating and deriving the consequences of alternative theories in such a way that the observations can actually be made to decide the question. Because two of the most important alternative theories are those which give rise to random distributions, methodology has come to be identified in the naive mind with statistical expertise. Statistical expertise is, of course, a particular kind of theoretical talent, a talent for deriving the implications of the theories which give rise to random distributions of one kind or another.

A strong esthetic reaction to crucial experiments is the central mark of the true scientist.

II / THE STRUCTURE OF CAUSAL THEORIES

Up to this point, we have not been concerned with the internal structure of our theory, called *A* above, but only with what happens to it as a whole when we conduct investigations to test it. In this section, we will analyze a particularly important class of theories, which are called *causal theories*. Not all scientific theories are causal theories in the sense in which we use the term here. By analyzing the internal logical structure of these theories, we can formulate more precisely how one derives empirical consequences from such theories. But before we can proceed with this analysis, we need to discuss the concept of a *variable* in scientific theories, because the *theoretical sentences*, or statements, of a causal theory are statements of a particular kind about the connection between scientific variables.

Definition of Variables

A "variable" in science is a *concept* which can have *various values*, and which is defined in such a way that *one can tell by*

means of observations which value it has in a particular occurrence. As will develop later, a causal law is a statement that certain values of two or more variables are connected in a certain way. The meaning of the terms of this definition can be better understood by considering a wide variety of variables used in sociological theory.

The simplest kind of variable is one which has *two values*, which can be represented as 1 and 0. Such a variable is commonly called a dichotomy, since it cuts the observations into two classes. There are first of all *natural* dichotomies, which we observe quite directly: sex, citizen or noncitizen of a particular country, employed or not employed, married or unmarried, all of which apply to people as the units of observation, and such things as votes in the United States (acts which can, usually, be classified as either Democratic or Republican), legal cases which can be classified as in federal or state jurisdiction, and so forth. The crucial point is that someone besides the investigator decides which of these classifications someone, or something, belongs to. The classifications exist independently of the scientific purposes of the investigator.

Then there are conceptual variables with two categories, in which the investigator himself creates the classification and places his observations in those two categories. For instance, in the comparative study of politics votes for different parties in different countries are often reduced to leftist votes and rightist votes; political scientists classify monarchies as constitutional or not constitutional, or governments as totalitarian or nontotalitarian.

Then there are variables with two values which are explicitly created as simplifications of variables which have more than two values. Some of this simplification may be done by the society itself, as in classifying people by age into minors and adults, or in simplifying the ranking system in the army into officers and men. Others are created by the investigator himself, as when an investigator has a scale of attitudes ranging from the extreme left to the extreme right and divides the population into leftists and rightists by splitting his scale in the middle, or when he has a scale of the degree of democracy of political systems and cuts it in the middle to distinguish dictatorships from democracies.

At the next level of complexity, we have variables which can take on a definite, known, finite number of values, which are gener-

ally called classifications. Some of these again are natural variables, as would be a classification of the population of the world by country of citizenship, ranks in the army, brands of automobiles or some other defined product, or votes classified by party in a country with more than two parties. Some variables with several values are conceptual, such as classification of building as single-family dwellings, apartment houses, commercial buildings, factory buildings, buildings with more than one of these functions, and "other buildings"; or a classification of governments into communist, liberal, and traditional; or a classification of workers according to their occupations or according to the industries in which they work. These variables can also be created by simplification of more complex variables, as when we divide incomes (a continuous variable) into high, medium, and low, or classify attitudes as rightist, leftist, and centrist. Finally, such variables can be created by combining two or more simpler variables into a complex variable. For instance, one could combine sex and marital status and obtain a variable with four values: married woman, married man, unmarried woman, unmarried man.

A third kind of variable is one which has exactly as many values as there are observations—that is, a complete ranking of the observations. An example would be rank in class at graduation.

Finally, there are continuous variables, or variables considered as continuous, such as the natural variables of income, age, floor area of a house, number of rooms in a house, years of employment in a firm; or such conceptual variables as intelligence quotient, degree of economic development of a country; or the constructed variables of the per cent in a given area voting for a candidate, or the size of a city.

For each of these types of variables we have a concept (*e.g.*, sex, income) in terms of which we make observations, and we classify or order these observations in some way (*e.g.*, into masculine and feminine, or into high, medium, and low income, or according to dollar income) so that each observation is connected with a single value of the variable.⁶

Sidney Siegel, *Non-Parametric Statistics for the Behavioral Sciences* (New York: McGraw-Hill, 1956), gives a good discussion of the logical character of different kinds of variables.

Definition of Causal Laws

Now we are in a position to describe what we mean by a causal law. A causal law is a statement or proposition in a theory which says that there exist environments (the better described the environments, the more complete the law) in which a change in the value of one variable is associated with a change in the value of another variable and can produce this change without any change in other variables in the environment.' It will be useful to use an example from another science to see more exactly what is involved in this definition.

Let us take "The shining of the sun causes the temperature to rise." Here we have two variables, whether the sun shines or not and temperature. We will call sun shining x and temperature y and write $x \rightarrow y$. The first is a dichotomy, a simplification of a continuous variable of the degree of sunshine; the second is a continuous variable. Then the causal law stated above means the following set of statements:

1. A change in x (in some defined environments) *is associated* with a change in y —there is a correlation between the two variables. When the sun shines, the temperature is on the average higher than when it does not shine.
2. One can produce the change in temperature by making the sun shine, but one cannot make the sun shine by changing the temperature.
3. There does not have to be any change in other variables for the sun to have its effect on the temperature (though, of course, the sun will have other effects than its effect on the temperature unless we control them, and the temperature itself will cause other variables to change unless they are controlled).

Here one should note the following important points:

1. There can be environments in which the law does not apply.

This definition is adapted from that of Russel Ackoff, *The Design of Social Research* (Chicago: University of Chicago Press, 1953), pp. 65-68. See also Herbert Simon, *Models of Man: Social and Rational* (New York: Wiley, 1957), chapter 3, pp. 50-61, and chapter 1, pp. 10-36.

It would not apply in a perfect vacuum, since temperature is not defined in such an environment.

2. The arrow with only one shaft, which we have used for causation, is quite different from the double-shafted arrow of logical implication, which we used above when discussing the *testing* of theories rather than their *formulation*.

3. The causal law can have variables of different classes, as dichotomies and continuous variables in this case.

4. There are other kinds of scientific theoretical propositions, than causal laws, such as systems of simultaneous equations without exogenous variables, or laws of mutual dependence without causal priority such as the relationship between temperature, pressure, and volume in a gas, or indirect causation ($x \rightarrow z \rightarrow y$), and so forth.

5. Variations in other variables can also cause variation in the caused variable, y , without falsifying the causal law. Any given dependent variable may be involved in a large number of causal laws.

6. Because of this, even though we know the causal law $x \rightarrow y$, we do not necessarily know that a *given* change in y that we observe is *in fact* caused by a variation in x , since it is quite possible that x does not vary in the environment we are investigating, or that the variations in y produced by variations in x are small relative to variations produced by other variables. For example, the sun can produce tides, but the major explanation for tides is the attraction of the moon. This means that even after a causal law is established, there is a further task of establishing which of the "natural" variations in y are *in fact* caused by x .

Observation in Support of Causal Theories

In order to derive observations sufficient to support or refute a causal theory, we must try to create observations of the following kinds:

1. We must observe *different values of the causal variable*. Unless there is variation in the causal variable, we cannot establish covariation. If one has a theory that individualism causes suicide, then one must observe at least two values of individualism, such as

Catholic and Protestant, educated and uneducated, commercial people and noncommercial, Jews and Gentiles. An observation of the suicide rates of Catholics alone, or of commercial people alone, is worthless for establishing covariation.

2. *Covariation*: We must observe variations in the dependent variable associated with these different values of the causal variable. Durkheim observed the suicide rates characteristic of groups with different degrees of individualism, thus classifying each group at the same time as high or low in individualism and as high or low in suicide rate.

3. *Causal Direction*: We must observe in some way that it is not possible to change the value of the causal variable by changing the supposed dependent variable, or that it *is* possible to change the value of the dependent variable by changing the causal variable.

4. *Nonspuriousness*: We must observe that there are not other variables in the environment which might cause changes in the dependent variable which change at the same time as the independent variable changes. There may be other effects of the causal variable which cannot be avoided, in which case one must try to show that they do not cause changes in the dependent variable. And there may also be effects of the dependent variable, in which case one also must try to establish the causal direction between the dependent variable and the possible confounding variable.

There are two main methods of observing *covariation*, the "experiment" and measurement of variables in their natural variations. The basic idea of the experiment is that the investigator himself changes the value of the causal variable for part of his observations, leaves it the same (or sets it at some other value) for another part, and measures the changes in the dependent variable.

In this way he has at least two values of the causal variable. These are often called the "treatment" and "control" observations, or the "experimental" and "control" observations. In the more advanced sciences usually the investigator sets several values of the causal variable to observe details of the variations of the dependent variable and has no "control" group. And he must measure variations in the dependent variable.⁸ The great advantage of this method

⁸Experimentation is more efficient if he can also measure variations in the causal variable.

of observing for covariation of two variables is that it solves at the same time the problem of causal direction, though not the problem of spuriousness.

The second method of observing for covariation is observation of natural variations of the two variables. In this case it is absolutely necessary that both variables be measured, and that there be natural variations of sufficient magnitude in the causal variable to have measurable effects. Sometimes such observations of natural variations are done with the help of special measuring instruments, such as interviews, tests, or calipers, and sometimes they are made without special aids (*e.g.*, by watching social interaction, or by the inspection of animals in zoology). Observing variations without special observational aids is often called natural history and is the preferred method of many anthropologists, political scientists, and historians in the social sciences. In any case, in order to establish covariation, one must observe at least two values of the independent variable (preferably many cases of each of the two) and measure or observe the associated variation of the dependent variable.

There seem to be five main methods for establishing *causal direction*: experimental manipulation of the causal variable, manipulation of the dependent variable (together with knowledge of covariation), temporal priority of changes or of determination of the value in one of the variables, knowing from other investigation the causes of observed variations in the causal variable, and knowing other causes of variation in the dependent variable if these causes are un-correlated with the variation in the independent variable. Each of these requires some explanation.

1. If we have manipulated the causal variable, then we know its cause in this particular case: our own action. If we know the cause of its variation (and if this cause cannot itself cause variations in the dependent variable), then we know that the covariation between the causal and the dependent variable must be due to the causal force of the causal variable.

2. Sometimes we cannot manipulate the causal variable—for instance, we may not be able to control the sunshine—but we can control the dependent variable—for instance, the temperature of a certain area. If we change the temperature of an area, and if the

sun does not come out or go in, then we know that the dependent variable does not cause the causal variable. If we know (through observing covariation) that there is some causal connection between the two variables, then the causal direction must go the other way around.

3. If we observe that the change in the causal variable (such as a rise in the price of cars) *precedes* changes in the dependent variable (such as an expansion in the production of cars), then we know that the second changes could not have produced the first change. Likewise if the value of some variable is determined, for some particular observations, at a time previous to observed differential changes of the dependent variable, then we know that the determination of the first variable caused (directly or indirectly) variations in the second variable. For instance, if we observe that children from higher-class families decide, during high school, to go to college with greater frequency than working-class children, we can establish causal direction if we have previously established that the class level of the family is mostly determined at the time the father enters the labor force.

4. If we know that the variation in the causal variable is due to other causes, as the variation of sunshine is due to the rotation of the earth on its axis, then we know that the variation in the temperature between day and night could not have caused the differences observed in the amount of sunlight between day and night. Hence we know that any observed covariation is due to sunshine causing variations in temperature.

5. If we know other causes of the dependent variable which are unrelated to the cause we are studying, then we can also establish causal direction. For all we need to do is to find observations in which the dependent variable varies (due to the other cause) without variation in the causal variable. For instance, differences in altitude cause systematic variations in temperature. If we choose an extremely dry climate, then the amount of sunshine will be equal on the heights and in the plains. If we observe covariation between altitude and temperature, and at given altitudes a covariation between sunshine and temperature, then we will have shown that some variations in temperature do not cause variations in the amount of sunshine, while variations of sunshine do cause variations in temperature.

It should be noted that some of these methods of establishing causal direction will show that the causal variable causes the variations in the dependent variable *even if* changes in the dependent variable can cause changes in the causal variable (especially the classical experiment, number 1 above, and sometimes number 3). Others will work only if there is no mutual causation (numbers 2, 4, and 5).

Finally, there are four main methods of establishing *nonspuriousness*: deliberate control of the value of possible spurious variables, control of possible spurious variables through randomization, control through knowing from the design of the investigation that the other variables will not vary between observations, and control through measurement of these variables and partialing out or correcting for their variation.

1. In the classical experiment as usually conceived in the physical sciences, other variables are deliberately controlled—for instance, the amount of impurities in a chemical reaction is deliberately set as near zero as possible, or temperature is controlled by air conditioning the laboratories. In order to do this, the variables which might cause variations in the dependent variable have to be pretty well known and they must be controllable through the deliberate intervention of the investigator. In general, such variables must be capable of being measured with high precision.

2. In control through randomization, one provides some sort of a *list* of the possible observed values of spurious variables and then chooses values from this list according to some random process for observation. For example, if we divide an experimental plot of land for a fertilization experiment into small plots, then we know that there is a value of natural fertility associated with each plot. Even though we cannot measure natural fertility (let us suppose), this provides us with a *list* of possible values of this spurious variable for various observations we can make. If we give each of these plots a number and choose which plots to apply artificial fertilizer to by using a table of random numbers, then we know that the values of natural fertility are controlled within certain statistical limits.⁹

⁹This is a rough way of talking about what is going on, and the degree of confidence we have in our conclusions will depend on the (unknown)

3. In control through knowing that all sets of observations have been exposed to the same values of third variables, even though we cannot control these variables, we ensure nonvariation of such variables and hence know that they could not have caused the variations in the dependent variable. For instance, in an agricultural experiment we may not be able to control the number of days of sunshine and rain to which fertilized and unfertilized plants are exposed, but we know enough about meteorology to know that if the plots are sufficiently close together they will be exposed to very small variations in weather.

4. Finally, we can measure third variables and compare covariation between our causal variable and the dependent variable only among observations where the third variable has identical values. Or if we know (or can compute from our data) the relation between the third variable and the dependent variable, we can correct the observed values of the dependent variable to take out the effect of the third variable and see whether there is still covariation between the causal variable and these corrected values. This is the *technique of partial correlation or standardization*.

In general, for any causal theory, then, one must derive empirical statements which specify observations which will establish co-variation, causal direction, and nonspuriousness. The particular practical situation of the investigator, and the nature of the variables, will determine which of the various kinds of observations he will be able to make and which will be most efficient. There are also various ways of combining the above techniques of observation to get greater efficiency. For instance, if one can measure some of the important spurious variables but not others but can provide some sort of list of different values of these others which he cannot measure, then *stratified random sampling* will improve the efficiency of the randomizing method of eliminating spuriousness.

This completes our logical analysis of causal theories and of how one derives empirical propositions from them which can be com-

shape of the distribution of natural fertility among plots. To speak more precisely, we will know with a certain probability that natural fertility differences would not account for differences of a certain size in the yield of the artificially fertilized plots as compared to the yield of unfertilized plots.

pared with observations. But in causal theories, as well as in other theories, *concepts* appear. Since these are keys to relating theories to empirical consequences, we must briefly discuss the logic of concept formation and its relation to empirical statements.

III / SCIENTIFIC CONCEPTS

Above we defined a scientific variable as a concept which can take on various values such that we can tell by observations which value it has in a particular case. Such direct observational concepts are among the most important in any science, though philosophers of science are generally inclined to admit that there are other concepts ("unobservables") in many or most scientific theories. Such concepts are electron, cause, a person's predisposition, and the like. We will deal here only with observational concepts, with emphasis on variables and types (complexes of variables).

Change of Concepts as Theories Change

The first requirement for a concept is that it accurately reflect the forces actually operating in the world. That is, the definition of a concept is a hypothesis that a certain sort of thing causes other things of interest to us. Usually this means that we have some specific ideas about what we want to explain, and that a certain kind of antecedent condition will in fact produce such phenomena as effects. If our theory is then refuted, we change the theory, which means among other things that we change our concepts or formulate new ones which more exactly correspond to the forces apparently operating.

For example, we may start with the idea that socialization in the lower class, in general, encourages juvenile delinquency. But on closer observation we find great variations in the amount of delinquency of children from family environments which seem to be, in the relevant respects, as nearly identical as we can measure. But we observe that some of the most delinquent children live in

certain neighborhoods and are far more concentrated than we would expect if only class factors were operating. We then form a concept of a delinquency-producing neighborhood and try to figure out how it might operate (so that we know what observations to make, what concept to form).

Suppose that on preliminary investigation of delinquency-producing neighborhoods, it seems that what is happening is that certain neighborhoods are places where delinquent teenage groups form and that individuals become delinquent mainly by learning from their teenage friends the values and practices of delinquency. So we move to a concept of a "delinquent subculture." When we have defined this sufficiently well and figured out how to measure it, we can return to the data to see whether working-class boys are any more likely to get into trouble than middle-class boys *when they have the same level of exposure to the delinquent subculture*.

But in order to decide how to conceptualize the delinquent subculture, we need to think again about how such a culture would work to produce delinquency. We might do this by specifying various values which we believe would motivate or justify delinquency, various skills which one might learn which, if he knew them, would make it more profitable to be a delinquent, and so forth. We would probably conduct investigations on the correlation of these values and these skills with delinquent behavior or with membership in known delinquent gangs. Once we had located a series of these values and skills, we would give a tentative definition of the delinquent subculture in terms of them.

But then we might notice that there were gangs with more or less the same values, with quite different rates of certain kinds of rational delinquency, according to the neighborhood they were located in (in other words, the delinquent subculture is not sufficient to explain all the variation in delinquency which varies with neighborhoods). Then we would be likely to redefine the concept of delinquency, to separate delinquencies caused by the delinquent subculture (perhaps vandalism, gang fighting, and so forth) and others which, though mostly occurring within it, require further explanation (rational crime). We might then form a concept of *opportunities* for rational crime made up of the existence of organized crime, markets for selling the numbers, and the like, in the

neighborhood. This part of delinquency then would be explained by delinquent subcultures plus opportunities."

As the science advances, it progressively redefines its concepts until they accurately represent the phenomena in the world. Both concepts defining the thing to be explained and the causal variables get redefined, until in the ideal case each concept represents phenomena which always have the same set of effects and the same set of causes, and all other characteristics of the observations are eliminated as irrelevant. Conceptual perfection cannot go on without the increase in knowledge about how the world works, for conceptual perfection is the location of phenomena with a unique set of causes and effects. It is quite useless to discuss concepts without reference to substantive theory about what goes on in the world, about what causes what. And such substantive theory is merely wind without observation ("research") to find out whether it is true or not.

Consequently every concept must be, either implicitly or explicitly, a *hypothesis* that specified phenomena, and no others, are, in some situations, causally operative. One does not formulate such concepts unless he has an idea that they cause something important or that they are caused by a distinct set of phenomena.

This means that concepts are in a constant state of flux as long as the causal theories are still in the process of development. And it means that the criteria for judging concepts are beliefs and evidence that the theories in which the concepts are involved are true. Usually when a theory proves inadequate, the concepts in it change (as, for instance, the concept of distance and time changed when relativity theory replaced Newtonian mechanics).

But because lower-level concepts are the part of a theory which directly corresponds to measurements or observations, they have some special characteristics. The practical and theoretical aspects of measurement and observation place certain requirements on the conceptual aspects of theories. We will discuss in turn the definition of *variables* and the definitions of *types*.

¹⁰ This represents roughly the conceptual development of one branch of the theory of juvenile delinquency. The names of Sutherland, Solomon Kobrin, Albert Cohen, Cloward, and Ohlin are associated with the main stages in this development.

Conceptualization of Variables

As we recall from the discussion above, a variable is a concept which has various values, such that one can tell from observations what value it has in a specific case. Usually there are several different ways of telling which value of a variable appears in a specific observation (which usually ought to agree, if they can be carried out at the same time). In fact, the correspondence is never exact, and measurement theory deals with approximations and criteria of agreement.

In general, a science starts off with its variables defined by common sense, by the distinctions that people make in daily life. Because people, in order to live efficiently, have to take account of the causal forces at work in the world, they make distinctions which are institutionalized in the language they speak. This is the level of "natural" variables which we discussed above, in which the investigator uses the values of variables given him by the society.

For instance, in discussing delinquency above, we implicitly started with the concept defined by the society, in which "delinquency" is activity of which the police and courts take account, and "juvenile delinquency" is delinquency committed by people whom, because of their age, the law treats in a special way. But by the time we got the theory developed, we were forced to redefine the concepts because different kinds of action that concern the police turned out to have different causes (distinguishing between rational crime and subculturally caused crime).

One of the fundamental difficulties with applied research generally is that natural variables *that create administrative problems* generally are not the same variables *that have a unique set of causes*. Sometimes applied researchers formulate this by saying that a natural variable "has multiple causes." From a scientific point of view, this means that the applied researcher is trying to explain the wrong thing.

In general, variables may be measured either by their *causes* or by their *effects*. Measurement of variables by their causes is most important in experimental research, where we try to manipulate the

independent variable, and measurement by effects is most important when we are measuring things as they occur naturally.

As an example of measurement by causes, a social psychologist might tell one group that "you will probably like each other" and another group that "you probably will not get along too well," in order to study the effects of social solidarity. Clearly he is measuring solidarity by its presumed causes (namely his statements to the subjects). On the other hand, intelligence presumably has the effect that a person is able to answer more questions of a certain kind on tests, and we use a series of these effects (a series of questions) to locate the underlying variable.

What this means, obviously, is that our measurement of any concept improves as our theories of its causes and effects improve. But there are certain techniques of improving measurement tools, without special theoretical analysis of the causal structure relating the underlying variable to those causes and effects by which we measure it.

These techniques depend on the idea that if two manifest observations are caused by the same underlying variable, there will be *covariation* between the observations, since there is causal covariation between the underlying variable and each of its effects. There will also be covariation between causes of an underlying variable and effects of that same variable. The techniques include factor analysis, Guttman scaling, latent structure analysis, item analysis, and so forth."

We will not go into these techniques here but will only point out that in the long run these are all implicitly causal theories of one kind or another, utilized for the special purpose of measurement of underlying variables. We use these causes and effects for measurement because we are not interested in these effects, or these causes, of the variables under investigation. One exercise a student can use to practice his theorizing is to take a group of related scales and guess at the causal structure involved. In factor analysis, this is called "interpretation" of the factors. Many people who use other techniques do not realize that their measuring instruments have an internal causal structure and hence do not "interpret" their results.

¹⁰ See Warren Torgerson, *Theory and Method of Scaling* (New York: Wiley, 1960).

The point here is that one uses the causes and effects of a variable to locate it, applying various techniques for analyzing the covariation between observations due to the observations' having common causes and common effects. This means that measurement is scientific theory in action for a specific purpose. The assumption of all these techniques is that if we can throw together enough minor causes and minor effects of an underlying variable (such as the minor effect of intelligence that it enables one to answer certain questions), those variables that really cause things and those that really are the result of a unique set of causes will turn up in our analysis. Then these are likely to be variables which cause the phenomena that we are really interested in explaining or which have a unique set of causes.

Further, the central way of increasing the covariation between observations that we use for measurement is to understand and conceptualize better the underlying causal structure. We can then obtain observations which are uniquely effects or causes of the variables actually operating and hence increase the covariation among observations. Thus the improvement of measurements is usually due to the advance of theory.

Measurement is not only a device for testing theory. It is a part of the theory.

Type-Concepts

A *type-concept* in scientific discourse is a concept which is constructed out of a *combination of the values of several variables*. Sometimes we find that in the world a whole series of variables has a set of values which are all the same in a large number of observations, and that if we find that one of the variables has a different value then all of them have different values. One group of type-concepts we are all familiar with is the chemical elements. If we isolate the elementary chemical substances and examine them, we observe that a large number of scientific variables such as valence, atomic weight, boiling and freezing points, specific gravity at a given temperature, number of atoms in a molecule, the strength of the bond they form in compounds, all go together. That is, all naturally occurring instances of elementary substances with an atomic weight

near 1 have the same valence, the same freezing and boiling points, the same specific gravity, the same number of atoms in a molecule, and the same strength of the chemical bond. We would call all these instances "hydrogen." If we look then at the instances of elementary substances with an atomic weight around 4, we observe a different valence, different boiling and freezing points, a different specific gravity, and a different strength of the chemical bond in compounds. We call all these instances "helium."

This fact that a large variety of variables takes on a limited number of *combinations of values* means that we simplify our theory greatly by talking about hydrogen and helium and the other elements rather than talking about all the values of all the different variables. Whenever a large number of variables go together, so that specific values of one are always associated with specific values of the others, the creation of *typologies*, or sets of type-concepts, such as the chemical elements, is scientifically useful. Other examples from various sciences are diseases in medicine (the variables are the symptoms which form the syndrome of the disease); compounds in chemistry; rock types in geology; the state-descriptions of solid, liquid, and gas in physics; classifications of societies as hunting-and-gathering, nomadic-herding, agricultural, and industrial in anthropology; classifications of languages according to their root language in linguistics; cloud-types in meteorology; and so forth.

The simplification of scientific theory by such typologies is due to the fact that many times the operative variable, either as cause or effect, is the type rather than the variables which make up the type. For a wide variety of chemical and physical problems we can formulate our predictions in terms of whether the gas we are working with is hydrogen or helium, forgetting about most of the variables which define the qualities of the two gases. Likewise we presume that the different diseases with different syndromes have different causes, while all instances of a given disease have the same causes.

We may, of course, find out that we have been wrong in such a supposition. In that case, we usually try to construct a new, better typology. When the causes of polio were discovered precisely, for instance, it turned out to be at least three different diseases.

The first test of a type-concept, then, is that the variables that

make it up be *in fact* connected to each other. We need to know that the variables have a combination of values in some instances and that these values are all the same in all such instances and have a different set of values in other instances. Thus in recent studies¹² evaluating Max Weber's type-concept of "bureaucracy," it has been discovered that part of the values of variables Weber used to construct his type-concept were in fact associated with each other, while another group of variables were not associated with these but were associated among themselves. Thus the concept had to be broken down into two different type-concepts ("rational" and "bureaucratic," or "professional" and "bureaucratic" administration).

The second criterion then is the criterion of all scientific concepts, that the typology be useful in the formulation of theories that are supported, that the type is indeed important as a cause or as an effect of other phenomena.

Types as a Convenience in Talking About Interaction Effects

There is another common use of typologies in scientific discourse which is not as fundamental as the simplification function—namely, to talk about *interaction effects* of two or more variables. By an "interaction effect" we mean that one variable has different effects, depending on the value some other variable has. For instance, people who are more interested in politics are more likely to attend political rallies for a candidate. But if people are very interested in politics but *favorable to the opposing candidate*, they are very unlikely to attend a political rally for the candidate. Thus variation in interest has different effects, depending on attitude toward the candidate.

It is often convenient in this situation, for simplicity in presentation, to create a new variable which takes account of both variables. That is, we define a new variable according to the *combinations*

¹²See the summary of studies by Stanley Udy and me in Peter Blau and Richard Scott, *Formal Organizations* (San Francisco: Chandler, 1962), especially pp. 207-08 on Stinchcombe and pp. 205-06 and 208-10 on Udy.

of values on other variables. Thus we might construct a variable in our example according to the following table:

INTEREST IN POLITICS	OPINIONS ON CANDIDATE A	
	<i>Favorable</i>	<i>Unfavorable</i>
High	Enthusiastic supporters	Enthusiastic opponents
Low	Apathetic followers	Apathetic opponents

Then if we were to relate this new variable to attendance at a political rally for candidate A, we would probably find that only the "enthusiastic supporters" were very likely to attend. That is, our typology constructed out of the two variables of interest and opinion would be more efficient than opinion alone or interest alone in predicting the attendance.

Such typologies for analyzing interaction effects are the most common typologies in sociology, and the fourfold table with types as entries in the cells is a standard tool of sociological theorizing. Notice that there is no statement in the table above that says that any combinations of values are more likely than any others. Nor is there anything comparable to the implicit statement in the periodic table of elements which says that elements with an atomic weight around 4 with a valence of +1 never occur. In typologies for analyzing interaction effects, all the combinations of values receive names and are presumed to be empirically possible.

There are, of course, other methods of handling interaction effects than the creation of typologies, since there are a large number of different ways of defining a new variable as a function of a combination of values on other variables. One of the most common ways of defining values of combinations with continuous variables is to form the product of the two. For instance, if we had a continuous measure of favorableness to candidate A and a continuous measure of political interest, then the product of these two variables would be high among the enthusiastic supporters and low in the other three cells, and thus would function in exactly the same way as our typology for predicting attendance at the rally.

Such typologies for handling interaction effects, then, are just one of a large class of ways of defining a new variable as a function of the combinations of values on other variables.

Examples of such nontypological combinations of variables to create a new variable would be the "cost of living" in a country (defined as a function of the prices of various goods), or "mean annual rainfall" (function of rainfall in various years). In each case the concept is useful because there are some effects of each of the variables which depend on the values of the others. The effect of a rise in the price of medical services on the welfare of workers, for instance, depends on whether other prices have gone up or down. A dry year will have different effects on what crops are planted the following spring if previous years have also been dry than if the area is a normally wet region. That is, there are interaction effects among the variables which make us want to use a function of their combinations as a theoretical variable rather than the variables themselves.

In summary, then, typologies have two radically different functions in scientific theory, one of which is fundamental, the other of which is just a convenience. In the first case a typology is a statement that a large number of variables have only a small number of combinations of values which actually occur, with all other combinations being rare or nonexistent. This results in a radical improvement in scientific theory. In the second case, a typology is merely a convenient way of writing a function of two or more variables in such a way that interaction effects can be simply stated.

IV / LEVELS OF GENERALITY IN SOCIAL THEORY

Up to this point we have discussed theories which are specific enough to have specific empirical consequences. We plan to spend much of our time at this level in this book, but it will be useful to have a general outline of all the various levels of generality which are included in the term "theory" as it is commonly used in sociology. Many of the debates and frictions among sociological researchers are really debates about the level of generality it is fruitful to work at. And many of the exchanges of criticisms,

though apparently about whether or not some particular theory is true, are actually conflicts over levels of generality.

Confusions over levels of generality quite often lead people to believe they have refuted something when they have not. To take a currently popular example, many people argue that since Marx predicted that the increasing misery of the workers would lead to their radicalization, and since this prediction turned out to be false, "Marxian theory" has been disproved. What has been disproved, of course, is only that part of Marxian theory which implied that prediction. Examining Marx's argument on this point, -we find it rather difficult to tell exactly what he does base this prediction on, but it seems to be his Ricardian analysis of the labor market. Thus there is certainly something wrong with the bald prediction, and very probably something wrong with Marx's version of Ricardo's theory of the price of labor. But it is quite possible that nothing is wrong with, for instance, his theory of politics as an expression of the class struggle, or with his theory that *if* there were increasing misery, *then* there would be a proletarian revolution. What has happened in this case (and it happens just as much to Freud as to Marx) is that people regard all elements of Marxian (or Freudian) theory as equally involved in every one of Marx's (or Freud's) hypotheses. Thus they refute, for instance, Marx's theory of politics by examining a consequence of his economic theory, or Freud's theory of the unconscious by refuting a specific theory of compulsive behavior. Many such problems can be avoided if we classify the elements of a man's thought according to the level of generality and use this distinction to guide our analysis of exactly what has happened, logically speaking, when a specific hypothesis of a theory has been refuted.

An Outline of Levels of Generality

It seems to me to be useful to classify elements of theories into the following seven levels of generality:

1. General ideas about causality, about what can be accepted as a fact, about what forms of logical inference are valid, and other similar philosophical presuppositions of scientific theories. This

chapter, for instance, is an example of writing at such a level. The general argument of Marx that the material world exists and all observable phenomena have material causes would be another example at this level.

2. General causal imageries, about the kinds of causes and causal structures that work for explaining phenomena of many varieties. Examples are: the idea of classical physics that causes cannot work from a distance without intermediaries; the sociological notion that some things are, or are not, to be explained by their consequences ("functional" theories); the idea that the physical environment is important in shaping animal and human behavior; or Marx's imagery that social relations in productive activities create interests and motives that people take to other areas of life.

3. Broad distinctions among classes of phenomena thought to have a distinctive type of explanation or to be the phenomena among which to search for causes of other phenomena. For instance, Freud distinguished between consciously controlled behavior and behavior that was not consciously monitored (slips of the tongue, dreams, hysterical symptoms, and reactions to projective techniques such as free associations) and held that the latter had a different set of causes from the former. Or Marx distinguished a set of authority relations and rights of appropriation ("property relations" or "relations of production") which he thought had a large number of effects and in turn were caused in a systematic way by the historical stage of development of the economy.

4. Ideas that the causes of one broad class of phenomena are likely to be found among variables in another broad class of phenomena. For instance, Freudian theory urges that the explanation of most behavior not controlled consciously is in the structure of unconscious motives, especially sexual, derived from repressions of these motives in infancy. Or Marx argued that most political phenomena are to be explained by variations in property relations.

5. Theories that one in particular of the variables within a broad class of phenomena explains a particular variable (or set of variables) in another class of phenomena. For instance, Marx argued that Bonapartism ("populist dictatorship," we would call it today)

was caused by a predominance of a "petty bourgeois" mode of production (one characterized by small businessmen and small farmers). He thought a petty bourgeoisie (especially one made up of small peasants) was both equalitarian and had great difficulty organizing as a class, needing therefore a democratically oriented dictator. Or Freud argued that hysterical paralysis was caused by a repression of a strong wish to do something involving the paralyzed member.

6. The empirical consequences of theories, describing the observations that could be made if the theory were true. Thus Marx argued that if petty bourgeois production causes Bonapartism, then

Louis Bonaparte in France would be heavily supported by petty bourgeois groups. A modern Marxist would argue in turn that the development of Bonapartism in Egypt or Mexico would have also been supported by small peasants and small businessmen. Or in Freud the empirical hypothesis would be that analysis of dreams, slips of the tongue, or free associations would reveal repressed wishes involving the paralyzed member in a particular case of hysterical paralysis.

7. Assertions that the observations in a particular case support, or refute, the empirical specification of level six. Thus Marx's assertions that in fact the petty bourgeois groups supported Louis Bonaparte, or Freud's assertion that in his case studies repressed wishes were indeed found, are of this order.

Levels of Critiques

A refutation or critique of a theory may touch any one of these levels. For the sake of illustration, let us go systematically through critiques at different levels, both for the Marxian theory of Bonapartism and for the Freudian theory of hysterical paralysis.

At the most general philosophical level, Lenin thought that Mach's philosophy of science attacked the materialist basis of Marxism, and that such petty bourgeois idealism led indirectly to considering the state as a mystical body, above mere class interest.

Skipping down to the fourth level, a critique might argue that the same mode of production tends to cause different political phenomena at different times in world history, depending on the

world ideological environment. For instance, radical movements are different after the 1917 Bolshevik Revolution than they were before. Thus the critique asserts that the causes of political forms are to be found in the reigning ideas of an age, rather than in the means of production. Marx would call such an argument also "idealistic," but it is clearly a different level of idealistic deviation from Marxism than Mach's.

An argument that populist dictatorships in the modern world (*e.g.*, in Egypt) are very often supported by the proletariat attacks one of Marx's theories at the fifth level. It does not necessarily challenge the fourth-level assumption that groups with a distinctive position in the productive system will have a distinctive politics. It merely says that Bonapartism in particular is one of the political expressions of the proletariat.

An argument that Marx incorrectly chose his indicators of a petty bourgeois productive structure, and that France in 1851 was really a feudal structure, would attack Marx at the sixth level. It says that the empirical assertion that France should have strong Bonapartist movements does not follow from the theoretical proposition that petty bourgeois modes of production ought to produce Bonapartist movements, for France was not in fact petty bourgeois. Likewise an argument that similar productive arrangements occur in the modern world and should produce Bonapartist movements is a modification at the sixth level. Countries with a large rural population and a land tenure system similar to that of France in the 1850's would include, in the 1960's: Haiti, Mexico, Formosa, Egypt, Ghana, Java in Indonesia, and perhaps India. An argument that these countries do not produce Bonapartist movements, yet should be included among countries with a petty bourgeois productive structure, would be an attack at the sixth level. This argument does not seem to me to show much promise.

An argument that Marx misinterpreted the facts about mid-nineteenth-century French political life and that it was really big business and big landowners who supported Louis Bonaparte would be an attack at the seventh level. The argument would agree that the facts ought to be the way Marx says they ought to be. But it says that in the real world the facts are some different way.

An argument that Freud had mistaken, say, a repressed fear for

a repressed sexual wish in one of his case studies would be an argument at the seventh level. It would be asserting that the facts are not as Freud says.

A derivation of new facts that Freud had not considered would be an attack (if the results came out negatively) at the sixth level. For instance, if Freud's theory were right, then it should be true that responses of hysterical paralysis patients to thematic apperception tests should differ from the responses of nonhysterical patients, by showing evidence of repressed sexually significant wishes using the paralyzed member. A critique that derived this consequence and then attacked the theory because this consequence is not true attacks at a higher level than a mere disputing of the facts.

The critique would immediately be moved up to the fifth level, if Freud were to accept this as a valid deduction from his theory. If the argument remains at the level of whether or not the thematic apperception results do derive from the theory, it need not move up.

An argument then might or might not develop at the fourth level. There are still other variables in the broad class of subconscious motives which might serve as alternatives to the repressed wish as explanations of hysterical symptoms. Proving that particular subconscious motives do not appear among hysterical patients on thematic apperception tests more frequently than among normals does not prove that some other motives may not appear more frequently. But one could assert that there are no such differences, nor any other evidence that subconscious motivations differ between hysterics and normals. Such an assertion of no subconscious differences would attack Freud on the fourth level.

An argument that both conscious and unconscious phenomena are learned in exactly the same way, according to the laws of learning psychology and reinforcement schedules, would be an attack at the third level. It would argue that exactly the same causes explain behavior which Freud classed in two distinct broad classes. It would argue that whether a piece of behavior was consciously monitored or not played no role in its explanation.

An argument that one cannot believe functional explanations of behavior, explanations by motives or other subjective orienta-

tions toward the ends of action, would be an attack at the second level. Behaviorism in psychology was such an attack.

An argument which says that since Freudian theory involves concepts which are unobservable (such as the ego and the id) it is all just a mystification would be an attack on fundamental philosophical postulates about what kinds of conceptual entities can enter into science.

It is extremely rare that a refutation at a lower level involves a refutation all the way back to the first principles of the philosophy of science. How much a theory is damaged by the refutation of a particular part always has to be analyzed logically."

V / CONCLUSIONS

In this chapter our purpose has been to formulate the logical tools needed to evaluate scientific theories, as a guide to the work of constructing these theories in such a form that they can be tested. In the first section we assumed that we could derive empirical statements from a theory we wanted to test. Then we analyzed what happened to our degree of belief in a theory when different kinds of consequences of the theory were compared with the facts. If a consequence of a theory turns out to be false, the

¹³ A good recent example of difficulties in this respect is George Homans and David Schneider, *Marriage, Authority, and Final Causes: A Study of Unilateral Cross Cousin Marriage* (Glencoe, Ill.: Free Press, 1955). They try to refute a particular "functional" theory that cross-cousin marriage is to be explained by its consequence that it ensures the regular and systematic exchange of women among families in certain kinds of kinship structures. Then they argue that they have shown that nothing is explained by its consequences. That is, after examining (and apparently refuting) a theory at level five (Rodney Needham argues that they in turn failed at level seven in assessing the facts of the case), they jumped up to level two to argue that this refutation showed the lack of usefulness of a general causal imagery. The general usefulness of a causal imagery is not, of course, involved in a failure of some particular theory which uses that causal imagery. We do not wish here either to judge whether their conclusions on functionalism are right (we will treat them later) or to evaluate their anthropological evidence. Our only point is that, if one accepts their evidence, they have only refuted a theory of the functions of cross-cousin marriage and have not shown that nothing is ever caused by its functions.

theory is falsified. If it turns out to be true, the theory becomes more credible. If several similar consequences turn out to be true, the theory becomes substantially more credible, for then more alternative theories are eliminated. The more different the consequences of a theory which are checked are, the more alternative theories are eliminated, and consequently checking many different consequences increases the credibility of a theory greatly.

But if an empirical consequence of a theory we are checking is also a consequence of either of the theories which generate random distributions (sampling or a large group of small causes), then the credibility of our theory is increased very little. Hence statistical significance of a set of observations (proof that they are unlikely to be explained by random distributions) is usually a minimum requirement for regarding the observations as substantial support for a theory.

If we can formulate the main alternatives to a theory, then we increase our scientific efficiency greatly by using crucial experiments. A crucial experiment is a set of observations which will give one result if one of the main alternative theories is true, and a different result if another is true.

Then we turned to the internal structure of one important class of scientific theories—namely, causal theories. In order to analyze such theories, we had first to define a scientific variable. A scientific variable is a concept which has several values. The concept and its values are defined in such a way that there is at least one way to tell what value an observation has in a particular instance (that is, at least one way to "measure" it). A causal law, then, is a statement that a change in the value of one variable is sufficient to produce a change in the value of another, without the operation of intermediate causes.

Observations in support of causal theories must support three elements: they must establish *covariation*, *causal direction*, and *nonspuriousness*. The two main ways of establishing covariation are experiment and measurement of simultaneous natural variation of two variables. The main methods of establishing causal direction are: manipulation of the independent variable, manipulation of the dependent variable together with knowledge of covariation, temporal priority of the causal variable, knowledge of the causes of the causal variable, and knowledge of some other causes

of the dependent variable which are not correlated with the cause we are studying. The main methods of establishing non-spuriousness are: deliberate control of possible spurious variables, control through randomization, control through knowledge that values of the possible spurious variable stay the same in various observations, and control through measurement of the spurious variable and either adjustment of the results for such spurious effects or partial analysis.

Then we turned to the next lower level in scientific theories, that of concepts. First we pointed out that the purpose of a concept is to identify phenomena which have identical causes or identical effects in some scientific field. Consequently concepts change as fast as the theories change, for as we locate the cause or the effects of phenomena more exactly, we must divide up phenomena differently. The test of a concept is how many supported scientific statements it occurs in.

The measurement of concepts is generally by either the causes of the phenomena or by their effects. The improvement of measurements and the refinement of concepts thus constitute a particular case of improvement of causal knowledge. However, the particular environment in which this investigation goes on has given rise to special devices for sorting out the causes and effects of a variable that are most directly related to the underlying variable, such as scaling techniques and factor analysis. All these devices depend on the fact that the direct effects or the direct causes of a variable we want to measure are likely to have high covariation with that variable, and hence high covariation with each other.

There are two main kinds of type-concepts, which use a combination of values on several variables to create a new variable. One of these, exemplified by the chemical elements, asserts that in the real world only a few of the many logically possible combinations of a number of variables actually occur; it gives names to each of these combinations. The other is a simple way of writing a new variable as a function of two or more others; it makes no assertions about the world. This last kind of type-concept is especially useful in talking about interaction effects of two or more variables, in which a change of a variable has a different effect depending on the values of other variables.

In the final section we turned to the general problem of how to

analyze more exactly what has happened to a whole theoretical structure when some particular derivation has been disproved. One important device for analyzing what happens is the classification of the elements of a theory according to their level of generality, ranging from general philosophical presuppositions down to assertions about what happened in a particular observation. A theoretical argument can be attacked at any one of these levels, but much confusion in scientific discourse has its origin in refuting a hypothesis at one level and believing that one has refuted a theoretical proposition at another level. The whole problem of the logical structure of complex theoretical systems has received very little fruitful analysis.

Although this chapter has touched on some problems currently discussed in the philosophy of science, we have tried to leave aside philosophical and epistemological problems whenever we could. Our purpose has not been to outline the ultimate justification for scientific belief, but to outline how scientific belief systems operate in practical fact, so that we can use this knowledge in constructing social theories.