CHAPTER 5

The Process of Stratification

Stratification systems may be characterized in various ways. Surely one of the most important has to do with the processes by which individuals become located, or locate themselves, in positions in the hierarchy comprising the system. At one extreme we can imagine that the circumstances of a person's birth—including the person's sex and the perfectly predictable sequence of age levels through which he is destined to pass—suffice to assign him unequivocally to a ranked status in a hierarchical system. At the opposite extreme his prospective adult status would be wholly problematic and contingent at the time of birth. Such status would become entirely determinate only as adulthood was reached, and solely as a consequence of his own actions taken freely—that is, in the absence of any constraint deriving from the circumstances of his birth or rearing. Such a pure achievement system is, of course, hypothetical, in much the same way that motion without friction is a purely hypothetical possibility in the physical world. Whenever the stratification system of any moderately large and complex society is described, it is seen to involve both ascriptive and achievement principles.

In a liberal democratic society we think of the more basic principle as being that of achievement. Some ascriptive features of the system may be regarded as vestiges of an earlier epoch, to be extirpated as rapidly as possible. Public policy may emphasize measures designed to enhance or to equalize opportunity—hopefully, to overcome ascriptive obstacles to the full exercise of the achievement principle.

The question of how far a society may realistically aspire to go in this direction is hotly debated, not only in the ideological arena but in the academic forum as well. Our contribution, if any, to the debate will consist largely in submitting measurements and estimates of the strength of ascriptive forces and of the scope of opportunities in a large contemporary society. The problem of the relative importance of the two principles in a given system is ultimately a quantitative one. We have pushed our ingenuity to its limit in seeking to contrive relevant quantifications.

The governing conceptual scheme in the analysis is quite a commonplace one. We think of the individual's life cycle as a sequence in time that can be described, however partially and crudely, by a set of classificatory or quantitative measurements taken at successive stages. Ideally we should like to have under observation a cohort of births, following the individuals who make up the cohort as they pass through life. As a practical matter we resorted to retrospective questions put to a representative sample of several adjacent cohorts so as to ascertain those facts about their life histories that we assumed were both relevant to our problem and accessible by this means of observation.

Given this scheme, the questions we are continually raising in one form or another are: how and to what degree do the circumstances of birth condition subsequent status? and, how does status attained (whether by ascription or achievement) at one stage of the life cycle affect the prospects for a subsequent stage? The questions are neither idle nor idiosyncratic ones. Current policy discussion and action come to a focus in a vaguely explicated notion of the "inheritance of poverty." Thus a spokesman for the Social Security Administration writes:

> It would be one thing if poverty hit at random and no one group were singled out. It is another thing to realize that some seem destined to poverty almost from birth—by their color or by the economic status or occupation of their parents.[1]

Another officially sanctioned concept is that of the "dropout," the person who fails to graduate from high school. Here the emphasis is not so much on circumstances operative at birth but on the presumed effect of early achievement on subsequent opportunities. Thus the "dropout" is seen as facing "a lifetime of uncertain employment,"[2] probable assignment to jobs of inferior status, reduced earning power, and vulnerability to various forms of social pathology.

1 Mollie Orshansky, "Children of the Poor," *Social Security Bulletin*, 26 (July 1963).

2 Forrest A. Bogan, "Employment of High School Graduates and Dropouts in 1964," *Special Labor Force Report*, No. 54 (U. S. Bureau of Labor Statistics, June 1965), p. 643.

In this study we do not have measurements on all the factors implicit in a full-blown conception of the "cycle of poverty" nor all those variables conceivably responding unfavorably to the achievement of "dropout" status. For practical reasons, as explained in Chapter 1, we were severely limited in the amount of information to be collected. For theoretical reasons—also spelled out more fully in Chapter 1—and in conformity with the tradition of studies in social mobility, we chose to emphasize occupation as a measure both of origin status and of status achievement. The present chapter is even more strictly limited to variables we think can be treated meaningfully as quantitative and therefore are suited to analysis by the regression technique described in Chapter 4. This limitation, however, is not merely an analytical convenience. We think of the selected quantitative variables as being sufficient to describe the major outlines of status changes in the life cycle of a cohort. Thus a study of the relationships among these variables leads to a formulation of a basic model of the process of stratification. In this chapter we consider also certain extensions of this model. Subsequent chapters provide, in effect, a number of additional detailed extensions, although these are secured only by giving up some of the elegance and convenience of the particular analytical procedures employed here.

A BASIC MODEL

To begin with, we examine only five variables. For expository convenience, when it is necessary to resort to symbols, we shall designate them by arbitrary letters but try to remind the reader from time to time of what the letters stand for. These variables are:

- V: Father's educational attainment
- X: Father's occupational status
- U: Respondent's educational attainment
- W: Status of respondent's first job
- Y: Status of respondent's occupation in 1962

Each of the three occupational statuses is scaled by the index described in Chapter 4, ranging from 0 to 96. The two education variables are scored on the following arbitrary scale of values ("rungs" on the "educational ladder") corresponding to specified numbers of years of formal schooling completed:

- 0: No school
- 1: Elementary, one to four years
- 2: Elementary, five to seven years

3: Elementary, eight years
4: High school, one to three years
5: High school, four years
6: College, one to three years
7: College, four years
8: College, five years or more (i.e., one or more years of postgraduate study)

Actually, this scoring system hardly differs from a simple linear transformation, or "coding," of the exact number of years of school completed. In retrospect, for reasons given in Chapter 4, we feel that the score implies too great a distance between intervals at the lower end of the scale; but the resultant distortion is minor in view of the very small proportions scored 0 or 1.

A basic assumption in our interpretation of regression statistics—though not in their calculation as such—has to do with the causal or temporal ordering of these variables. In terms of the father's career we should naturally assume precedence of *V* (education) with respect to X (occupation when his son was 16 years old). We are not concerned with the father's career, however, but only with his statuses that comprised a configuration of background circumstances or origin conditions for the cohorts of sons who were respondents in the OCG study. Hence we generally make no assumption as to the priority of *V* with respect to X; in effect, we assume the measurements on these variables to be contemporaneous from the son's viewpoint. The respondent's education, *U*, is supposed to follow in time—and thus to be susceptible to causal influence from—the two measures of father's status. Because we ascertained X as of respondent's age 16, it is true that some respondents may have completed school before the age to which X pertains. Such cases were doubtlessly a small minority and in only a minor proportion of them could the father (or other family head) have changed status radically in the two or three years before the respondent reached 16.

The next step in the sequence is more problematic. We assume that *W* (first job status) follows *U* (education). The assumption conforms to the wording of the questionnaire (see Appendix B), which stipulated "the first full-time job you had after you left school." In the years since the OCG study was designed we have been made aware of a fact that should have been considered more carefully in the design. Many students leave school more or less definitively, only to return, perhaps to a different school, some years later, whereupon they often

finish a degree program.[3] The OCG questionnaire contained information relevant to this problem, namely the item on age at first job. Through an oversight no tabulations of this item were made for the present study. Tables prepared for another study[4] using the OCG data, however, suggest that approximately one-eighth of the respondents report a combination of age at first job and education that would be very improbable unless (a) they violated instructions by reporting a part-time or school-vacation job as the first job, or (b) they did, in fact, interrupt their schooling to enter regular employment. (These "inconsistent" responses include men giving 19 as their age at first job and college graduation or more as their education; 17 or 18 with some college or more; 14, 15, or 16 with high-school graduation or more; and under 14 with some high school or more.) When the two variables are studied in combination with occupation of first job, a very clear effect is evident. Men with a given amount of education beginning their first jobs early held lower occupational statuses than those beginning at a normal or advanced age for the specified amount of education.

Despite the strong probability that the *U-W* sequence is reversed for an appreciable minority of respondents, we have hardly any alternative to the assumption made here. If the bulk of the men who interrupted schooling to take their first jobs were among those ultimately securing relatively advanced education, then our variable *W* is downwardly biased, no doubt, as a measure of their occupational status immediately after they finally left school for good. In this sense, the correlations between *U* and *W* and between *W* and *Y* are probably attenuated. Thus, if we had really measured "job after completing education" instead of "first job," the former would in all likelihood have loomed somewhat larger as a variable intervening between education and 1962 occupational status. We do not wish to argue that our respondents erred in their reports on first job. We are inclined to conclude that their reports were realistic enough, and that it was our assumption about the meaning of the responses that proved to be fallible.

The fundamental difficulty here is conceptual. If we insist on *any* uniform sequence of the events involved in accomplishing the transi-

[3] Bruce K. Eckland, "College Dropouts Who Came Back," *Harvard Educational Review*, 34(1964), 402-420.

[4] Beverly Duncan, *Family Factors and School Dropout: 1920–1960*, U. S. Office of Education, Cooperative Research Project No. 2258, Ann Arbor: Univers. of Michigan, 1965.

tion to independent adult status, we do violence to reality. Completion of schooling, departure from the parental home, entry into the labor market, and contracting of a first marriage are crucial steps in this transition, which all normally occur within a few short years. Yet they occur at no fixed ages nor in any fixed order. As soon as we aggregate individual data for analytical purposes we are forced into the use of simplifying assumptions. Our assumption here is, in effect, that "first job" has a uniform significance for all men in terms of its temporal relationship to educational preparation and subsequent work experience. If this assumption is not strictly correct, we doubt that it could be improved by substituting any other *single* measure of initial occupational status. (In designing the OCG questionnaire, the alternative of "job at the time of first marriage" was entertained briefly but dropped for the reason, among others, that unmarried men would be excluded thereby.)

One other problem with the *U-W* transition should be mentioned. Among the younger men in the study, 20 to 24 years old, are many who have yet to finish their schooling or to take up their first jobs or both—not to mention the men in this age group missed by the survey on account of their military service (see Appendix C). Unfortunately, an early decision on tabulation plans resulted in the inclusion of the 20 to 24 group with the older men in aggregate tables for men 20 to 64 years old. We have ascertained that this results in only minor distortions by comparing a variety of data for men 20 to 64 and for those 25 to 64 years of age. Once over the *U-W* hurdle, we see no serious objection to our assumption that both *U* and *W* precede *Y*, except in regard to some fraction of the very young men just mentioned.

In summary, then, we take the somewhat idealized assumption of temporal order to represent an order of priority in a causal or processual sequence, which may be stated diagrammatically as follows:

$$(V, X) - (U) - (W) - (Y).$$

In proposing this sequence we do not overlook the possibility of what Carlsson calls "delayed effects,"[5] meaning that an early variable may affect a later one not only via intervening variables but also directly (or perhaps through variables not measured in the study).

In translating this conceptual framework into quantitative estimates the first task is to establish the pattern of associations between the variables in the sequence. This is accomplished with the correlation coefficient, as explained in Chapter 4. Table 5.1 supplies the correla-

[5] Gösta Carlsson, *Social Mobility and Class Structure*, Lund: CWK Gleerup, 1958, p. 124.

TABLE 5.1. SIMPLE CORRELATIONS FOR FIVE STATUS VARIABLES

Variable	Variable: Y	W	U	X	V
Y: 1962 occ. status	...	.541	.596	.405	.322
W: First-job status		...	.538	.417	.332
U: Education			...	.438	.453
X: Father's occ. status				...	.516
V: Father's education					...

tion matrix on which much of the subsequent analysis is based. In discussing causal interpretations of these correlations, we shall have to be clear about the distinction between two points of view. On the one hand, the simple correlation—given our assumption as to direction of causation—measures the gross magnitude of the effect of the antecedent upon the consequent variable. Thus, if $r_{YW} = .541$, we can say that an increment of one standard deviation in first job status produces (whether directly or indirectly) an increment of just over half of one standard deviation in 1962 occupational status. From another point of view we are more concerned with net effects. If both first job and 1962 status have a common antecedent cause—say, father's occupation—we may want to state what part of the effect of *W* on *Y* consists in a transmission of the prior influence of *X*. Or, thinking of *X* as the initial cause, we may focus on the extent to which its influence on *Y* is transmitted by way of its prior influence on *W*.

We may, then, devote a few remarks to the pattern of gross effects before presenting the apparatus that yields estimates of net direct and indirect effects. Since we do not require a causal ordering of father's education with respect to his occupation, we may be content simply to note that $r_{XV} = .516$ is somewhat lower than the corresponding correlation, $r_{YU} = .596$, observed for the respondents themselves. The difference suggests a heightening of the effect of education on occupational status between the fathers' and the sons' generations. Before stressing this interpretation, however, we must remember that the measurements of *V* and *X* do not pertain to some actual cohort of men, here designated "fathers." Each "father" is represented in the data in proportion to the number of his sons who were 20 to 64 years old in March 1962.

The first recorded status of the son himself is education (*U*). We note that r_{UV} is just slightly greater than r_{UX}. Apparently both measures on the father represent factors that may influence the son's education.

In terms of gross effects there is a clear ordering of influences on first job. Thus $r_{WU} > r_{WX} > r_{WV}$. Education is most strongly corre-

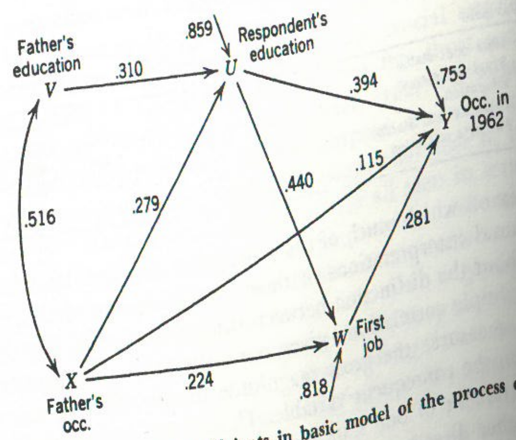

Figure 5.1. Path coefficients in basic model of the process of stratification.

lated with first job, followed by father's occupation, and then by father's education.

Occupational status in 1962 (Y) apparently is influenced more strongly by education than by first job; but our earlier discussion of the first-job measure suggests we should not overemphasize the difference between r_{YW} and r_{YU}. Each, however, is substantially greater than r_{YX}, which in turn is rather more impressive than r_{YV}.

Figure 5.1 is a graphic representation of the system of relationships among the five variables that we propose as our basic model. The numbers entered on the diagram, with the exception of r_{XV}, are path coefficients, the estimation of which will be explained shortly. First we must become familiar with the conventions followed in constructing this kind of diagram. The link between V and X is shown as a curved line with an arrowhead at both ends. This is to distinguish it from the other lines, which are taken to be paths of influence. In the case of V and X we may suspect an influence running from the former to the latter. But if the diagram is logical for the respondent's generation, we should have to assume that for the fathers, likewise, education and occupation are correlated not only because one affects the other but also because common causes lie behind both, which we have not measured. The bidirectional arrow merely serves to sum up all sources of correlation between V and X and to indicate that the explanation thereof is not part of the problem at hand.

The straight lines running from one measured variable to another represent *direct* (or net) influences. The symbol for the path coeffi-

cient, such as p_{YW}, carries a double subscript. The first subscript is the variable at the head of the path, or the effect; the second is the causal variable. (This resembles the convention for regression coefficients, where the first subscript refers to the "dependent" variable, the second to the "independent" variable.)

Finally, we see lines with no source indicated carrying arrows to each of the effect variables. These represent the residual paths, standing for all other influences on the variable in question, including causes not recognized or measured, errors of measurement, and departures of the true relationships from additivity and linearity, properties that are assumed throughout the analysis (as explained in the section on regression in Chapter 4).

An important feature of this kind of causal scheme is that variables recognized as effects of certain antecedent factors may, in turn, serve as causes for subsequent variables. For example, U is caused by V and X, but it in turn influences W and Y. The algebraic representation of the scheme is a system of equations, rather than the single equation more often employed in multiple regression analysis. This feature permits a flexible conceptualization of the *modus operandi* of the causal network. Note that Y is shown here as being influenced directly by W, U, and X, but not by V (an assumption that will be justified shortly). But this does not imply that V has no influence on Y. V affects U, which does affect Y both directly and indirectly (via W). Moreover, V is correlated with X, and thus shares in the gross effect of X on Y, which is partly direct and partly indirect. Hence the gross effect of V on Y, previously described in terms of the correlation r_{YV}, is here interpreted as being entirely indirect, in consequence of V's effect on intervening variables and its correlation with another cause of Y.

PATH COEFFICIENTS

Whether a path diagram, or the causal scheme it represents, is adequate depends on both theoretical and empirical considerations. At a minimum, before constructing the diagram we must know, or be willing to assume, a causal ordering of the observed variables (hence the lengthy discussion of this matter earlier in this chapter). This information is external or *a priori* with respect to the data, which merely describe associations or correlations. Moreover, the causal scheme must be complete, in the sense that all causes are accounted for. Here, as in most problems involving analysis of observational data, we achieve a formal completeness of the scheme by representing unmeasured causes as a residual factor, presumed to be uncorrelated with the remaining factors lying behind the variable in question. If any factor is known or presumed to operate in some other way it must be represented in the diagram in accordance with its causal role, even though it is not measured. Sometimes it is possible to deduce interesting implications from the inclusion of such a variable and to secure useful estimates of certain paths in the absence of measurements on it, but this is not always so. A partial exception to the rule that all causes must be explicitly represented in the diagram is the unmeasured variable that can be assumed to operate strictly as an intervening variable. Its inclusion would enrich our understanding of a causal system without invalidating the causal scheme that omits it. Sociologists have only recently begun to appreciate how stringent are the logical requirements that must be met if discussion of causal processes is to go beyond mere impressionism and vague verbal formulations.[6] We are a long way from being able to make causal inferences with confidence, and schemes of the kind presented here had best be regarded as crude first approximations to adequate causal models.

On the empirical side, a minimum test of the adequacy of a causal diagram is whether it satisfactorily accounts for the observed correlations among the measured variables. In making such a test we employ the fundamental theorem in path analysis, which shows how to obtain the correlation between any two variables in the system, given the path coefficients and correlations entered on the diagram.[7] Without stating this theorem in general form we may illustrate its application here. For example,

$$r_{YX} = p_{YX} + p_{YU}r_{UX} + p_{YW}r_{WX};$$

and

$$r_{WX} = p_{WX} + p_{WU}r_{UX}.$$

We make use of each path leading to a given variable (such as Y in the first example) and the correlations of each of its causes with all other variables in the system. The latter correlations, in turn, may be analyzed; for example, r_{WX}, which appeared as such in the first equation, is broken down into two parts in the second. A complete expansion along these lines is required to trace out all the indirect connections between variables; thus,

$$r_{YX} = p_{YX} + p_{YU}p_{UX} + p_{YU}p_{UV}r_{VX} + p_{YW}p_{WX} + p_{YW}p_{WU}p_{UX} + p_{YW}p_{WU}p_{UV}r_{VX}.$$

6 H. M. Blalock, Jr., *Causal Inferences in Nonexperimental Research*, Chapel Hill: Univer. of North Carolina Press, 1964.

7 Sewall Wright, "Path Coefficients and Path Regressions," *Biometrics*, 16 (1960), 189–202; Otis Dudley Duncan, "Path Analysis," *American Journal of Sociology*, 72(1966), 1-16.

Now, if the path coefficients are properly estimated, and if there is no inconsistency in the diagram, the correlations calculated by a formula like the foregoing must equal the observed correlations. Let us compare the values computed from such a formula with the corresponding observed correlations:

$$\begin{aligned} r_{WV} &= p_{WX}r_{XV} + p_{WU}r_{UV} \\ &= (.224)(.516) + (.440)(.453) \\ &= .116 + .199 = .315 \end{aligned}$$

which compares with the observed value of .332; and

$$\begin{aligned} r_{YV} &= p_{YU}r_{UV} + p_{YX}r_{XV} + p_{YW}r_{WV} \\ &= (.394)(.453) + (.115)(.516) + (.281)(.315) = .326 \end{aligned}$$

(using here the calculated rather than the observed value of r_{WV}), which resembles the actual value, .322. Other such comparisons—for r_{YX}, for example—reveal, at most, trivial discrepancies (no larger than .001).

We arrive, by this roundabout journey, at the problem of getting numerical values for the path coefficients in the first place. This involves using equations of the foregoing type inversely. We have illustrated how to obtain correlations if the path coefficients are known, but in the typical empirical problem we know the correlations (or at least some of them) and have to estimate the paths. For a diagram of the type of Figure 5.1 the solution involves equations of the same form as those of linear multiple regression, except that we work with a recursive system of regression equations[8] rather than a single regression equation.

Table 5.2 records the results of the regression calculations. It can be seen that some alternative combinations of independent variables were studied. It turned out that the net regressions of both W and Y on V were so small as to be negligible. Hence V could be disregarded as a direct influence on these variables without loss of information. The net regression of Y on X was likewise small but, as it appears, not entirely negligible. Curiously, this net regression is of the same order of magnitude as the proportion of occupational inheritance in this population—about 10 per cent, as discussed in Chapter 4. We might speculate that the direct effect of father's occupation on the occupational status of a mature man consists of this modest amount of strict occupational inheritance. The remainder of the effect of X on Y is indirect, inasmuch as X has previously influenced U and W, the son's education and the occupational level at which he got his start. For reasons noted in Chapter 3 we do not assume that the full impact of

8 Blalock, *op. cit.*, pp. 54ff.

TABLE 5.2. PARTIAL REGRESSION COEFFICIENTS IN STANDARD FORM (BETA COEFFICIENTS) AND COEFFICIENTS OF DETERMINATION, FOR SPECIFIED COMBINATIONS OF VARIABLES

Dependent Variable[a]	Independent Variables[a] W	U	X	V	Coefficient of Determination (R^2)
U[b]	...	...	.279	.310	.26
W	...	.433	.214	.026	.33
W[b]	...	.440	.224	...	.33
Y	.282	.397	.120	-.014	.43
Y[b]	.281	.394	.115	...	.43
Y	.311	.428	...	...	.42

[a]V: Father's education.
X: Father's occ. status.
U: Respondent's education.
W: First-job status.
Y: 1962 occ. status.
[b]Beta coefficients in these sets taken as estimates of path coefficients for Figure 5.1.

the tendency to take up the father's occupation is registered in the choice of first job.

With the formal properties of the model in mind we may turn to some general problems confronting this kind of interpretation of our results. One of the first impressions gained from Figure 5.1 is that the largest path coefficients in the diagram are those for residual factors, that is, variables not measured. The residual path is merely a convenient representation of the extent to which measured causes in the system fail to account for the variation in the effect variables. (The residual is obtained from the coefficient of determination; if $R^2_{Y(WUX)}$ is the squared multiple correlation of Y on the three independent variables, then the residual for Y is $\sqrt{1 - R^2_{Y(WUX)}}$.) Sociologists are often disappointed in the size of the residual, assuming that this is a measure of their success in "explaining" the phenomenon under study. They seldom reflect on what it would mean to live in a society where nearly perfect explanation of the dependent variable could be secured by studying causal variables like father's occupation or respondent's education. In such a society it would indeed be true that some are "destined to poverty almost from birth . . . by the economic status or occupation of their parents" (in the words of the reference cited in footnote 1). Others, of course, would be "destined" to affluence or to modest circumstances. By no effort of their own could they materially alter the course of destiny, nor could any stroke of fortune, good or ill, lead to an outcome not already in the cards.

Thinking of the residual as an index of the adequacy of an explanation gives rise to a serious misconception. It is thought that a high multiple correlation is presumptive evidence that an explanation is correct or nearly so, whereas a low percentage of determination means

that a causal interpretation is almost certainly wrong. The fact is that the size of the residual (or, if one prefers, the proportion of variation "explained") is *no* guide whatever to the validity of a causal interpretation. The best-known cases of "spurious correlation"—a correlation leading to an egregiously wrong interpretation—are those in which the coefficient of determination is quite high.

The relevant question about the residual is not really its size at all, but whether the unobserved factors it stands for are properly represented as being uncorrelated with the measured antecedent variables. We shall entertain subsequently some conjectures about unmeasured variables that clearly are not uncorrelated with the causes depicted in Figure 5.1. It turns out that these require us to acknowledge certain possible modifications of the diagram, whereas other features of it remain more or less intact. A delicate question in this regard is that of the burden of proof. It is all too easy to make a formidable list of unmeasured variables that someone has alleged to be crucial to the process under study. But the mere existence of such variables is already acknowledged by the very presence of the residual. It would seem to be part of the task of the critic to *show,* if only hypothetically, but *specifically,* how the modification of the causal scheme to include a new variable would disrupt or alter the relationships in the original diagram. His argument to this effect could then be examined for plausibility and his evidence, if any, studied in terms of the empirical possibilities it suggests.

Our supposition is that the scheme in Figure 5.1 is most easily subject to modification by introducing additional measures of the same kind as those used here. If indexes relating to socioeconomic background other than V and X are inserted we will almost certainly estimate differently the direct effects of these particular variables. If occupational statuses of the respondent intervening between W and Y were known we should have to modify more or less radically the right-hand portion of the diagram, as will be shown in the next section. Yet we should argue that such modifications may amount to an enrichment or extension of the basic model rather than an invalidation of it. The same may be said of other variables that function as intervening causes. In theory, it should be possible to specify these in some detail, and a major part of the research worker's task is properly defined as an attempt at such specification. In the course of such work, to be sure, there is always the possibility of a discovery that would require a fundamental reformulation, making the present model obsolete. Discarding the model would be a cost gladly paid for the prize of such a discovery.

Postponing the confrontation with an altered model, the one at hand is not lacking in interest. An instructive exercise is to compare the magnitudes of gross and net relationships. Here we make use of the fact that the correlation coefficient and the path coefficient have the same dimensionality. The correlation $r_{YX} = .405$ (Table 5.1) means that a unit change (one standard deviation) in X produces a change of 0.4 unit in Y, in gross terms. The path coefficient, $p_{YX} = .115$ (Figure 5.1), tells us that about one-fourth of this gross effect is a result of the direct influence of X on Y. (We speculated above on the role of occupational inheritance in this connection.) The remainder ($.405 - .115 = .29$) is indirect, via U and W. The sum of all indirect effects, therefore, is given by the difference between the simple correlation and the path coefficient connecting two variables. We note that the indirect effects on Y are generally substantial, relative to the direct. Even the variable temporally closest (we assume) to Y has "indirect effects"—actually, common antecedent causes—nearly as large as the direct. Thus $r_{YW} = .541$ and $p_{YW} = .281$, so that the aggregate of "indirect effects" is .26, which in this case are common determinants of Y and W that spuriously inflate the correlation between them.

To ascertain the indirect effects along a given chain of causation we must multiply the path coefficients along the chain. The procedure is to locate on the diagram the dependent variable of interest, and then trace back along the paths linking it to its immediate and remote causes. In such a tracing we may reverse direction once but only once, following the rule "first back, then forward." Any bidirectional correlation may be traced in either direction. If the diagram contains more than one such correlation, however, only one may be used in a given compound path. In tracing the indirect connections no variable may be intersected more than once in one compound path. Having traced all such possible compound paths, we obtain the entirety of indirect effects as their sum.

Let us consider the example of effects of education on first job, U on W. The gross or total effect is $r_{WU} = .538$. The direct path is $p_{WU} = .440$. There are two indirect connections or compound paths: from W back to X then forward to U; and from W back to X, then back to V, and then forward to U. Hence we have:

$$\underset{\text{(gross)}}{r_{WU}} = \underset{\text{(direct)}}{p_{WU}} + \underbrace{p_{WX}p_{UX} + p_{WX}r_{XV}p_{UV}}_{\text{(indirect)}}$$

or, numerically,

$$\begin{aligned} .538 &= .440 + (.224)(.279) + (.224)(.516)(.310) \\ &= .440 + .062 + .036 \\ &= .440 + .098. \end{aligned}$$

In this case all the indirect effect of U on W derives from the fact that both U and W have X (plus V) as a common cause. In other instances, when more than one common cause is involved and these causes are themselves interrelated, the complexity is too great to permit a succinct verbal summary.

A final stipulation about the scheme had best be stated, though it is implicit in all the previous discussion. The form of the model itself, but most particularly the numerical estimates accompanying it, are submitted as valid only for the population under study. No claim is made that an equally cogent account of the process of stratification in another society could be rendered in terms of this scheme. For other populations, or even for subpopulations within the United States, the magnitudes would almost certainly be different, although we have some basis for supposing them to have been fairly constant over the last few decades in this country. The technique of path analysis is not a method for discovering causal laws but a procedure for giving a quantitative interpretation to the manifestations of a known or assumed causal system as it operates in a particular population. When the same interpretive structure is appropriate for two or more populations there is something to be learned by comparing their respective path coefficients and correlation patterns. We have not yet reached the stage at which such comparative study of stratification systems is feasible.

AGE GROUPS: THE LIFE CYCLE OF A SYNTHETIC COHORT

For simplicity, the preceding analysis has ignored differences among age groups. Our present task is to venture some interpretation of such differences. The raw material for the analysis is presented in Table 5.3 in the form of simple correlations between pairs of the five status variables under study. For the reasons mentioned in Chapter 3, this analysis is confined to men with nonfarm background.

We must consider immediately what kinds of inferences or interpretations are allowed by comparisons among the four cohorts. Three of the variables are specified as of a more or less uniform stage of the respondent's life cycle: father's occupation (X), respondent's education (U), and first job (W). Father's education (V), on the other hand, was presumably determinate in the father's youth; the time interval between V and any of the former variables would be determined in large part by father's age at respondent's birth. This interval is variable in length. We might, however, assume that the time interval from V to X, though highly variable within each cohort of respondents, has a similar average and dispersion from one cohort to another. If father's education is taken as a fixed status once the father has completed his