1 Background

This book is intended for graduate students and postgraduate researchers in social and personality psychology who wish to build on a foundation of graduate-level training in analysis of variance and multiple regression analysis and familiarity with factor analysis to learn the basics of structural equation modeling.[1] The book offers a nontechnical overview at a level of depth designed to prepare readers to read and evaluate reports of research, and begin planning their own research, using structural equation modeling. This concise, application-oriented treatment is no substitute for coursework and fuller written treatments such as can be found in textbooks (e.g., Bollen, 1989; Kaplan, 2009; Kline, 2010; Schumacker & Lomax, 2004). Rather, it provides a bridge between briefer treatments offered in book chapters and didactic journal articles (e.g., Hoyle, 2007; Weston & Gore, 2006) and those fuller treatments.

Structural equation modeling (SEM) is a term used to describe a growing and increasingly general set of statistical methods for modeling data. In light of the statistical training typical of researchers in social and personality psychology, two features of SEM are worth noting at the outset. First, unlike familiar statistical methods such as analysis of variance (ANOVA) and multiple regression analysis, which estimate parameters (e.g., means, regression coefficients) from individual cases, SEM estimates parameters from variances and covariances. Indeed, it is possible to apply most forms of SEM using only the variances of a set of variables and the covariances between them. Second, although this feature is not required and, on occasion, is not used, a significant strength of SEM is the capacity to model relations between latent variables; that is, the unobserved constructs of which observed variables are fallible representations. The focus on covariances rather than data from individual cases involves a move away from familiar estimators such as ordinary least squares toward more general estimators such as maximum likelihood. And the capacity to model relations between latent variables shifts the focus of data analysis from variables to constructs, thereby more closely aligning conceptual and statistical expressions of hypotheses. These departures from the statistical methods traditionally used by social and personality researchers require thinking differently about how data are brought to bear on research questions and hypotheses. The payoff is a comprehensive approach to modeling data that is well suited for empirical tests of the richly detailed, process-oriented models of the human experience typical of social and personality psychology.

In this opening chapter, I lay the groundwork for a nontechnical presentation of the nuts and bolts of SEM in the next three chapters and for promising applications in social and personality psychology in the final chapter. I begin by positioning SEM among the array of statistical methods of which most researchers in social and personality psychology would be aware. I then provide a sketch of the relatively short history of SEM. I next offer working definitions of key concepts with which many readers are not familiar. I conclude the chapter with a section on the use of path diagrams to convey a model to be estimated using SEM.

SEM in Statistical Context

One way to make clear the comprehensiveness of SEM and the flexibility it affords for modeling data is to compare it to the various statistical methods historically used in social and personality research. Before making these comparisons, it bears noting that the full set of techniques captured by the term "structural equation modeling" is ever expanding, so that the term now invokes a substantially broader set of techniques than when it came into standard usage in the late 1970s. Indeed, the continuing expansion of SEM capabilities has made the boundaries between SEM and other statistical approaches somewhat hard to define. With that caveat in mind, I can describe SEM in relation to traditional and emerging statistical methods.

The names of statistical methods commonly used by researchers in social and personality psychology and selected newer methods are displayed in Figure 1.1. The methods are arrayed from specific to general moving from left to right. An arrow from one method to the next indicates that the former is a limited form of the latter. Put differently, methods to which arrows point could be used to address any statistical hypothesis that could be addressed by methods from which those arrows originate and any methods linked to them. For instance, referring to the top line in the figure, *t*-test is a limited form of ANOVA because it allows for a single factor with no more than two levels, whereas ANOVA accommodates multiple factors with two or more levels. ANOVA is the specific instance of ordinary least squares (OLS) regression in which all predictors are categorical. And so on.

We first encounter a form of SEM, *covariance structure modeling*, about midway across the figure. As evidenced by the arrows pointing to it, this elemental form of SEM incorporates the essential capabilities of multiple regression analysis and factor analysis. What is not apparent from the figure (and not essential information for the present discussion) is that covariance structure modeling is far more than simply a hybrid of these two well-known statistical strategies. Nevertheless, a useful starting point for researchers in social and personality

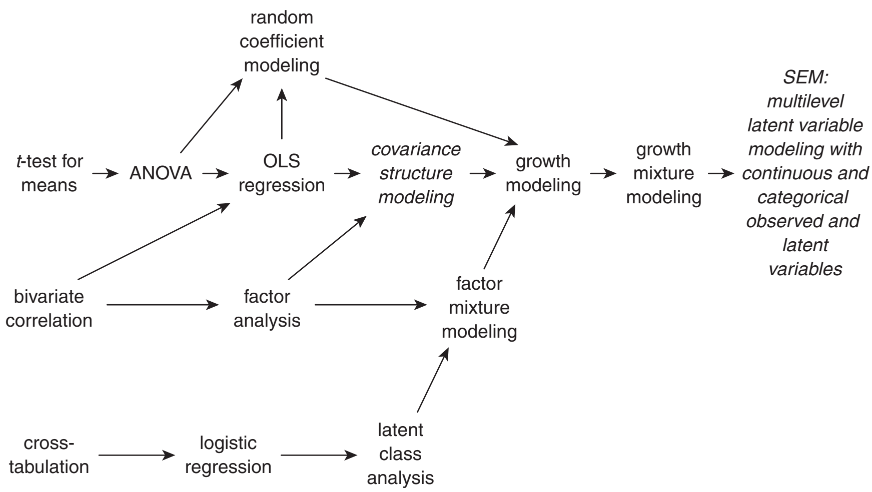

Figure 1.1 Selected statistical methods in relation to structural equation modeling

psychology is the realization that, in simple terms, SEM can be conceptualized as an extension of multiple regression analysis in which multiple equations (often with multiple, directionally related outcomes) are simultaneously estimated and both predictors and outcomes can be modeled as the equivalent of factors as traditionally modeled in separate analyses using factor analysis.

A limitation of this form of SEM is the focus solely on covariances between variables that are assumed to be measured on continuous scales. Examination of other arrows leading to and away from the covariance structure modeling entry in the figure makes clear how SEM is expanded by incorporating means modeling and allowing for categorical variables. The addition of means to variances and covariances in the data matrix allows for the modeling of patterns of means over time in *latent growth models* (described in Chapter 5) as well as other models analyzed using *random coefficient modeling*. Analyses made possible by the ability to estimate parameters from categorical data are traced from specific to general along the bottom of Figure 1.1. The transition from *logistic regression* (and other forms of the generalized linear model) to *latent class analysis* represents a shift from categories manifest in the observed data to those that are unobserved and inferred from data. Traditional *factor analysis* and *latent class analysis* meet in *factor mixture modeling*, which allows for both continuous and categorical latent variables. The addition of means gives rise to *growth mixture modeling*, in which heterogeneity in trajectories of means is modeled as latent classes. The array culminates on the right in the long-winded but apt label, *multilevel latent variable modeling with continuous and categorical observed and latent variables*, which reflects the current breadth of models that can be estimated within the SEM framework.

A somewhat arbitrary, primarily semantic, but nonetheless instructive distinction that can be made between the methods shown in Figure 1.1 is between those used to *analyze* data and those used to *model* data (Rodgers, 2010). Of course, ANOVA, the prototypic analysis method, can be accurately described as a strategy for modeling sources of variance in individual data; however, ANOVA is typically used in such a way that a relatively standard model is virtually always applied (main effects and interactions). Custom contrasts, trend analyses and the like reflect a move from simply analyzing data from individual cases to modeling means. Factor analysis and latent class analysis reflect the transition from analyzing to modeling data. They suffer, however, from the limitation that, as typically used, models are discovered rather than posited and tested. From covariance structure modeling forward, the methods are best described as approaches to modeling data rather than simply analyzing data. This distinction is elaborated further in the next chapter.

A further distinction is between methods that typically are used to analyze or model individual cases and those typically used to model relations between variables. I noted at the outset that SEM differs from statistical methods commonly used by researchers in social and personality psychology in its focus on covariances rather than data from individual cases. To elaborate further, parameter estimation and statistical tests in familiar methods such as ANOVA and multiple regression analysis typically are based on the principle of minimizing the sum of the squared differences between observed scores for individual cases on the dependent variable and the case-level scores predicted by the independent variables. "Error" is defined as the average squared difference between observed and predicted scores across all cases. The goal of estimation in SEM is the minimization of the difference between the observed covariance matrix and the covariance matrix implied by the model. "Error" is defined as the difference between these two matrices as reflected in the value of an estimator-specific fit function (covered in Chapter 3). In both cases the focus is on the degree to which a model either prescribed by the typical application of the method (as in ANOVA and multiple regression analysis) or specified by the researcher (as in SEM) reproduces the observed data. The distinction is in what constitutes the observed data – case-level scores in ANOVA and multiple regression analysis, variances, and covariances in SEM.

I close this section by noting a final pair of observations inspired by Figure 1.1. As established earlier, a statistical method can be used to accomplish any statistical hypothesis test that could be accomplished using methods prior to it in the figure. For instance, two means, which might typically be compared using the *t*-test, also could be compared using ANOVA, multiple regression analysis, and SEM. Although the use of SEM to compare two means could be defended on statistical grounds, practically speaking it would be unwise. The *t*-test is perfectly suited to this hypothesis test and requires little explanation or justification.

Following this principle of using the simplest and most straightforward statistical method for hypothesis tests, SEM becomes relevant when the hypothesis is a model that implies multiple equations (i.e., statements of the relations between independent and dependent variables) and/or makes use of a data set that includes multiple indicators of key constructs, allowing for the expression of constructs as latent variables.

Although SEM is not always recommended for hypothesis testing in social and personality research, knowledge of the full array of modeling possibilities offered by SEM can inspire predictions and models that might not otherwise have been ventured. Research questions and explanatory models, on occasion, lead to the development of statistical methods for addressing them (e.g., models of intelligence and factor analysis). Typically, however, questions and models are shaped by the statistical methods of which the researcher is aware. As such, the more flexible and general the statistical approach, the broader the range of research questions and explanatory models likely to be ventured. The range of modeling possibilities afforded by SEM suggest new ways for social and personality psychologists to think about social and personality processes, pose research questions, design research, and analyze data.

Historical Roots

Most historical accounts of the development of SEM trace its origins to the early 1920s and the development of path analysis by Sewall Wright, an evolutionary biologist. Wright invented the statistical method of *path analysis*, a graphical model in which the linear relations between variables are expressed in terms of coefficients that are derived from the correlations between them (Wright, 1934). The potential value of Wright's model for social research was not immediately recognized; it was not until the 1960s that applications of path analysis to social research data were more fully developed. The principal figures in early applications of path analysis to data from social research were sociologists Blalock (1964) and Duncan (1966). Duncan and Goldberger, an econometrician, integrated the sociological approach to path analysis with the simultaneous equations approach in economics (e.g., Goldberger & Duncan, 1973) and the factor analytic approach in psychology (e.g., Duncan, 1975; Goldberger, 1971), yielding the basic form of SEM.

This general model was formalized and extended in the 1970s by Jöreskog (1973), Keesling (1972), and Wiley (1973), producing what became known as the LISREL (*L*inear *S*tructural *REL*ations) model. This model includes two parts, one specifying the relations between indicators and latent variables – the *measurement model* – and the other specifying the relations between latent variables – the *structural model*. The LISREL model served as the basis for the LISREL software program, which, by the release of Version 3 in the mid-1970s, allowed substantively oriented social and behavioral researchers to specify, estimate, and test latent variable models using SEM.

The earliest uses of SEM in social and personality psychology appeared in the late 1970s and early 1980s. The earliest published applications in social psychology were by Peter Bentler and colleagues. For example, Bentler and Speckart (1979) used SEM to model the relation between attitude and behavior expressed as latent variables, including the first statistical modeling of the full set of relations in the influential theory of reasoned action (Fishbein & Ajzen, 1975). The earliest published uses of SEM in personality research are more difficult to pinpoint; however, by the mid-1980s applications of SEM, particularly the measurement model, to questions regarding personality structure began to appear (e.g., Reuman, Alwin, & Veroff, 1984; Tanaka & Bentler, 1983; Tanaka & Huba, 1984). In a prototypic application, Reuman et al. used SEM to model the achievement motive as a latent variable and show that, when random measurement error is separated from reliable common variance in fallible measures of the construct, validity coefficients are consistent with the theoretical model of the construct.

By the late 1980s, spurred by compelling applications such as those by Bentler and Speckart (1979) and Reuman et al. (1984), and better access to software for specifying and estimating models, SEM found traction in social and personality psychology. Its use, and particularly the interpretation of its results, quickly gave rise to a literature on misuses of SEM and misinterpretation of SEM results by psychological scientists (e.g., Baumrind, 1983; Breckler, 1990; Cliff, 1983). The use of SEM in social and personality psychology has improved and increased despite the fact that formal training in SEM in social and personality psychology doctoral programs is still not the norm (Aiken, West, & Millsap, 2008). Extracurricular workshops and didactic volumes such as this one are the means by which many researchers in social and personality psychology learn about the capabilities of SEM and acquire basic proficiency in its use (often as implemented in a specific statistical software package). Although SEM is not likely to join ANOVA and multiple regression analysis as statistical methods that all social and personality researchers must know, its use will no doubt increase as compelling applications are published with increasing frequency.

The Language of SEM

Terminology

As with any statistical method (though perhaps more so), SEM is characterized by terminology that takes on precise and specific meaning when used with reference to it. Key terms are given full treatment at appropriate points later in the book. Basic definitions, which are offered in this section, will allow me to use the terms selectively to provide an initial sketch of SEM in the remainder of this chapter and the first part of Chapter 2.

Perhaps the most basic term in the SEM lexicon is *model*, a formal statement about the statistical relations between variables. Models typically are conveyed in diagrams as shown later in this chapter, or as equations as shown in Chapter 2. The origin and evaluation of models in the SEM context vary according to the modeling approach taken (Jöreskog, 1993). In the *strictly confirmatory* approach, the goal is to evaluate the degree to which a single, a priori model accounts for a set of observed relations between variables. For example, a researcher might evaluate the degree to which self-ratings on a set of adjectives selected to represent four types of affect conform to a model in which each adjective loads on only one of four correlated factors. Alternatively, instead of focusing on a single model, SEM might be used to compare two or more competing models in the *alternative models* approach. To continue the example, in addition to a model with four correlated factors, the researcher might consider a model with four uncorrelated factors, a model with a single factor, and/or a second-order model in which covariation between the four factors is explained by one superordinate factor. Finally, the use of SEM might involve *model generating*. If, for example, the researcher's proposed four-factor model does not adequately explain self-ratings on the adjectives, and there are no obvious alternative models, rather than abandoning the data, the researcher might use them to generate a model. Of course, using the data to generate a model of the data is a questionable practice (MacCallum, Roznowski, & Necowitz, 1992); however, careful modification of an a priori model with the goal of finding a model that accounts for the data can lead to tentative discoveries that ultimately result in amended or revised theoretical models

Specification involves designating the variables, relations between variables, and status of the parameters in a model. In terms of *designating variables*, the decisions are which variables in a data matrix to include as measured variables and which latent variables, if any, to model. In terms of *designating the relations between variables*, the researcher must decide which variables are related and, for those that are related, whether the relation is nondirectional or directional. Finally, the *status of parameters* in a model must be specified. Although, strictly speaking, specification is always involved in tests of statistical hypotheses, it is, in most cases, accomplished without the knowledge of the researcher in social or personality psychology. For example, the standard model estimated in applications of ANOVA – all main effects and interactions – typically is specified without consideration for other models that might be fit to the data. In typical applications of multiple regression analysis a specification decision is required in order to position one variable as the outcome and the others as predictors. Perhaps the closest that researchers in social and personality psychology come to specification is in decisions required for hierarchical multiple regression (e.g., how many sets; which variables in which sets; order in which sets enter?) and exploratory factor analysis (e.g., number of factors to extract; method of rotation?). Because there is no standard model to be fitted using SEM, any application requires specification. Detailed coverage is provided in Chapter 2.

A key aspect of specification is designating the status of parameters (e.g., variances, covariances, factor loadings, regression coefficients) in a model. Although specification can be quite specific regarding both the magnitude and sign of parameters, parameters typically are specified as either fixed or free. *Fixed parameters* are not estimated from the data and their value typically is fixed at zero or 1.0. *Free parameters* are estimated from the data. Because data analytic methods traditionally used in social and personality research do not focus on modeling, readers might not be aware of the fixed and free parameters in applications of those methods. In a hierarchical multiple regression analysis that includes three sets of variables, each comprising two variables, Step 1, which appears to include only two variables, could alternatively be viewed as a model that includes all six variables in which the regression weights for variables in the second and third sets have been fixed at zero. At Step 2, when variables in the second set are added, two of the four formerly fixed parameters (i.e., regression weights) are now free. In the alternative models approach described earlier, the differences between models to be compared typically involve parameters that are free in one model and fixed in the other. In the model generating approach, the adjustment of an initial model in an attempt to better account for the data often involves freeing parameters that were fixed and, to a lesser extent, fixing parameters that were free.

The parameters of most interest in models are those associated with *paths*, which signify directional relations between two variables as in the effect of a predictor on an outcome in multiple regression analysis. The *path coefficient* indicates the magnitude and strength of the effect of one variable on another. Virtually all models include *direct effects*, which propose that one variable is temporally or causally antecedent to one other variable. These are types of effects routinely estimated in ANOVA or multiple regression analysis. Within a model, each direct effect characterizes the relation between an independent and a dependent variable, though the dependent variable in one direct effect can be the independent variable in another. Moreover, like multiple regression, a dependent variable can be related to multiple independent variables, and, like multivariate analysis of variance, an independent variable can be related to multiple dependent variables. The capacity to treat a single variable as both a dependent and independent variable lies at the heart of the *indirect effect*, the effect of an independent variable on a dependent variable through one or more intervening, or mediating, variables (Baron & Kenny, 1986). In the case of a single mediating variable, the mediating variable is a dependent variable with reference to the independent variable but an independent variable with reference to the dependent variable. Thus, the simplest indirect effect involves two direct effects. For instance, if x has a direct effect on z, and z has a direct effect on y, then x is said to have an indirect effect on y through z. The sum of the direct and indirect effects of an independent variable on a dependent variable is termed the *total effect* of the independent variable.

Effects in models involve one or both of two types of variables. *Observed variables* (sometimes referred to as manifest variables) are those for which there are values in the case-level data matrix. Virtually all analytic methods traditionally used by researchers in social and personality psychology estimate effects between observed variables. SEM also allows for the estimation of effects involving *latent variables*, which are implied by a model but are not represented by values in the case-level data matrix. Latent variables, or factors, are a function of observed variables, which, when used to model a latent variable, are referred to as *indicators*. Indicators are of two types, formative and reflective (Cohen, Cohen, Teresi, Marchi, & Velez, 1990). *Formative indicators* are presumed to cause their latent variable, which is modeled as a weighted, linear combination of them as in principal components analysis. *Reflective indicators* are presumed to be caused by their latent variable, which is modeled as an explanation of the commonality among them as in principal factors analysis. Although latent variables can, in principle, be a function of formative indictors, the overwhelming majority of latent variables are a function of the commonality among a set of reflective indicators as with common factors in exploratory factor analysis (Edwards & Bagozzi, 2000; more on this distinction in Chapter 5). Variables, whether observed or latent, can be further distinguished according to whether they are exogenous or endogenous. *Exogenous variables* (i.e., independent variables) are those for which no explanation is attempted within the model; that is, there are no paths directed toward them. *Endogenous variables* (i.e., dependent variables) are those to which one or more paths are directed within the model.

It is virtually always the case that some portion of the variance in endogenous variables is not explained by paths directed toward them. In such cases, unexplained variance is described in one of two ways depending on how the variable is positioned in the model. In the case of latent variables for which indicators are reflective, the unexplained variance in indicators is attributed to *uniquenesses*, variance unique to a given indicator in the sense that it is not shared with other indicators of the latent variable. In the case of latent or observed (nonindicator) endogenous variables, variance not accounted for by variables in the model that predict them is allocated to *disturbances* (equivalent to the error term in a regression equation). As will soon be apparent, uniquenesses and disturbances are latent variables that in some models can be specified as related to other variables in a model.

A fully specified model, with its observed and latent variables and fixed and free parameters, implies a structure that is not directly evident in the unstructured set of $p(p - 1)/2$ covariances (where p is the number of observed variables; this value does not include variances). Although much can be learned from a thorough examination of the covariances (particularly in standardized form as correlation coefficients), the degree to which they are consistent with theory-based models that offer accounts of the relations between the variables can rarely be determined from the data in their most elemental form. A specified model proposes a *structure* or pattern of statistical relations that is more useful, interesting, and parsimonious than the bivariate associations in the covariance matrix (hence the occasional reference to SEM as covariance structure analysis). As described below and detailed in Chapter 3, the question of model fit can be expressed as how well the covariance structure offered by the model maps onto the unstructured set of covariances.

Although it would seem that, research design and logical considerations aside, any arrangement of variables and set of relations between them is possible with SEM, such is not the case. A key consideration when specifying a model is *identification*, which concerns whether a single, unique value for each and every free parameter can be obtained from the observed data. If for each free parameter a unique estimate can be obtained through one and only one manipulation of the observed data, then the model is *just identified* and has zero degrees of freedom. If a unique estimate for one or more free parameters can be obtained in multiple ways from the observed data, then the model is *overidentified* and has degrees of freedom equal to the number of observed variances and covariances minus the number of free parameters. If a single, unique estimate cannot be obtained from the observed data for one or more free parameters, then the model is *underidentified* and cannot be validly estimated. Thus, a restriction on specification is that the resultant model must be either just identified or overidentified. Although identification is rarely a concern in statistical models traditionally used by social and personality researchers, researchers occasionally stumble on it as a result of the inadvertent inclusion of a continuous variable as a factor in an ANOVA, resulting in a model requiring more degrees of freedom than the $N - 1$ that are available (i.e., it is underidentified). Identification is covered in greater detail in Chapter 2.

A properly specified model can be estimated. *Estimation* is the statistical process of obtaining estimates of free parameters from observed data. Although single-stage least squares methods such as those used in standard ANOVA or multiple regression analysis can be used to derive parameter estimates, iterative methods such as maximum likelihood or generalized least squares are preferred. *Iterative methods* involve a series of attempts to obtain estimates of free parameters that imply a covariance matrix like the observed one. The *implied covariance matrix* is the covariance matrix that would result if values of fixed parameters and estimates of free parameters were substituted into the structural equations, which then were used to derive a covariance matrix. Iteration begins with a set of *start values*, tentative values of free parameters from which an implied covariance matrix can be computed and compared to the observed covariance matrix. After each iteration, the resultant implied covariance matrix is compared to the observed matrix. The comparison between the implied and observed covariance matrices results in a residual matrix. The *residual matrix* contains elements whose values are the differences between corresponding values in the implied and observed matrices. Iteration continues until it is not possible to update the parameter estimates and produce an implied covariance matrix whose elements are any closer in magnitude and direction to the corresponding elements in the observed covariance matrix. Said differently, iteration continues until the values of the elements in the residual matrix cannot be minimized any further. At this point the estimation procedure is said to have *converged*.

A properly converged solution produces the raw materials from which various statistics and indices of fit are constructed. A model is said to *fit* the observed data to the extent that the covariance matrix it implies is equivalent to the observed covariance matrix (i.e., elements of the residual matrix are near zero). The question of fit is, of course, a statistical one that must take into account features of the data, the model, and the estimation method. For instance, the observed covariance matrix is treated as a population covariance matrix, yet that matrix suffers from sampling error – increasingly so as sample size decreases. Also, the more free parameters in a model, the more likely the model is to fit the data because parameter estimates are derived from the data. Chapter 3 reviews several statistics and indices of fit, highlighting how each accounts for sampling error and lack of parsimony.

As described at the beginning of this section, one way in which SEM can be applied is the alternative models approach, which involves comparing models that offer competing accounts of the data. Such models cannot always be formally compared. In some instances two or more alternatives are *equivalent models*; that is, they produce precisely the same implied covariance matrix and, as a result, identical fit to the data. Ideally, two or more models to be compared are not only not equivalent – they are nested. Two models are *nested* if they both contain the same parameters but the set of free parameters in one model is a subset of the free parameters in the other. Such models can be formally compared and, on statistical grounds, one chosen over the other. Readers familiar with hierarchical linear regression, in which predictors are entered in sets and statistical significance judged by the R^2 increment, already understand the general idea of nested models.

One possible outcome of the strictly confirmatory and alternative models approaches to SEM is that the model(s) posited a priori do not provide an acceptable account of the data. In such cases, the researcher can either abandon the analysis or move to a model generating approach. This move entails *model modification* (or respecification), freeing parameters that in the a priori model(s) were fixed and/or fixing parameters that were free. Such decisions are made through *specification searching*, which may involve either a diagnosis by the researcher based on evaluation of output (e.g., the residual matrix) or an automated search based on statistical criteria implemented by the software. This exercise may produce a model that appears to offer an acceptable account of the data, but such models always await confirmation using new data.

Although not exhaustive, the list of terms defined here is sufficient to begin an exploration of SEM, with the goal of describing and illustrating applications well suited to research in social and personality psychology. This coverage of foundational information about SEM concludes with an overview of path diagrams.

Path Diagrams

The models typically tested using methods such as ANOVA and multiple regression analysis have become somewhat standardized. Moreover, the models are straightforward, involving a single dependent variable and a set of independent variables for which linear, and sometimes interactive, effects are estimated. As such, relatively little description or explanation of how these methods are being applied in a given study is required. As will become clear in Chapter 2, there is no standard application of SEM and, for a given set of variables, a potentially large number of models could be specified. This state of affairs makes it necessary for researchers to accurately and completely describe the model(s) to be fitted. An effective means of conveying the details of a model is the path diagram.

In all likelihood the path diagram originated with Sewall Wright, who, as mentioned earlier, developed path analysis. The earliest instances appeared in a 1920 article on the relative contribution of heredity and environment to color variation in guinea pigs, which also introduced the terms *path* and *path coefficient*. The first instance, although it includes the directional arrows commonplace in path diagrams, also includes sketches of a guinea pig dam and sire as well as two offspring that vary in coloration. Moreover, the symbols for the genetic contributions to the color of offspring are a sperm and an egg! A second figure is both less entertaining and remarkably similar to path diagrams routinely used by sociologists in the 1960s and 1970s.

As path analysis has been subsumed by SEM and SEM has expanded, the demands on path diagrams as a means of conveying the details of a model have increased. Indeed, some models are sufficiently complex that the path diagram is no longer an effective means of communicating the details of the model. For most models that would be specified by researchers in social and personality psychology, however, the path diagram is an effective and efficient means of describing models to be estimated using SEM.

A path diagram that includes most of the elements typical of path diagrams is shown in Figure 1.2. The general flow of the diagram is from left to right. In this instance, the model begins with *F1*, which is presumed to arise outside the model (i.e., it is exogenous), and culminates with *F3*, the construct the model is presumed to explain or, statistically speaking, account for variance in. When possible, the constructs are arrayed according to their presumed position in the model. In this instance, *F2* is set between *F1* and *F3* because, as will soon be apparent, it is hypothesized to mediate the relation between *F1* and *F3*.

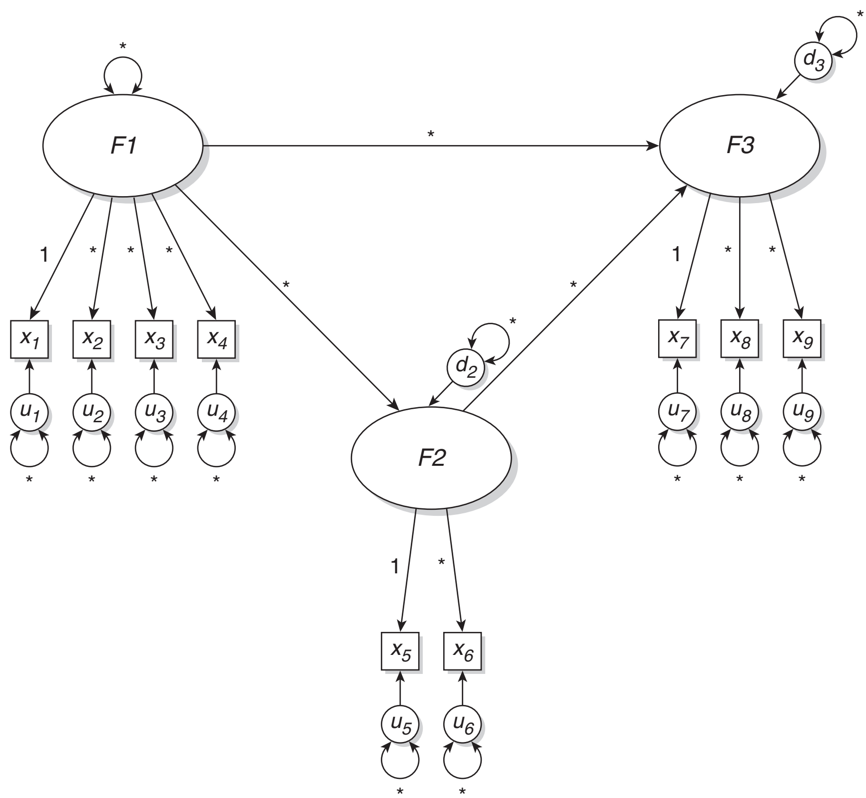

Figure 1.2 Example of a path diagram

With this general orientation in mind, let us now consider the components of the path diagram. The ovals and circles represent latent variables, sources of influence not measured directly. The ovals correspond to substantive latent variables, or factors. The oval labeled *F1* is an independent variable – it is not influenced by other variables in the model. The ovals labeled *F2* and *F3* are dependent variables – their variance is, in part, accounted for by other variables in the model. Paths run from each of these latent variables to their indicators, represented by squares labeled x_1 to x_9. These paths are either labeled "1," which means the factor loading has been fixed at this value (rationale provided in Chapter 2), or "*," indicating that the factor loading is free and must be estimated from the data. Variance in each indicator not attributable to the latent variable is allocated to measurement error, or uniqueness, indicated by the small circles labeled u_1 to u_9. Associated with each of these circles is a curved, two-headed arrow and a *, which indicates a variance. The three latent variables are connected by directional arrows. Associated with each is a path coefficient, accompanied by a * indicating the magnitude and direction of influence of one latent variable on another. Small circles also are associated with the endogenous latent variables. These indicate disturbances, variance in the latent variables, labeled d_2 and d_3, not accounted for by other latent variables in the model. Finally, there is a variance, indicated by *, associated with the latent independent variable.

As is true of most models, this model includes a combination of free and fixed parameters. Free parameters are easily identified by the *. The location of fixed parameters is less obvious. It is apparent that there is a single fixed loading on each latent variable. The remaining fixed loadings involve paths that could have been included but were not. For instance, there is no path from *F1* to x_5. Implicitly, this path has been fixed to zero. Also, there are no covariances between uniquenesses, meaning these parameters are implicitly fixed at zero as well. Fixed parameters in the form of excluded paths are desirable in a model, for they contribute to parsimony and overidentification. They also can explain the inadequacy of a poor-fitting model. Hence, when processing path diagrams, it is important to take note of paths that have been omitted, indicating that the accompanying parameters have been fixed at zero.

One additional feature of the model in Figure 1.2 bears mention. Notice that the directional effect of *F1* on *F3* takes two forms in the model. The path diagram indicates that *F1* has a direct effect on *F3* as indicated by the horizontal path along the top of the diagram. In addition, the model indicates that *F1* has an indirect effect on *F3* through *F2*. That is, *F2* serves as an intervening variable, or mediator, through which the effect of *F1* on *F3* is transmitted.

Modeling the variables of primary interest, *F1*, *F2*, and *F3*, as latent variables takes advantage of a key strength of SEM over traditional statistical approaches in social and personality psychology – the capacity to model out unreliability, thereby producing estimates of directional relations that have been corrected for attenuation (Muchinsky, 1996). The advantage of this capacity is aptly illustrated in the model depicted in Figure 1.2. The more unreliable the indicators of the intervening variable, *F2*, the greater the underestimation of the effect of *F2* on *F3* and overestimation of the effect of *F1* on *F2* (Hoyle & Kenny, 1999) if *F2* is not modeled as a latent variable. In other words, one might fail to support the prediction that *F2* mediates the *F1–F3* relation when, in reality, it does (more on this in Chapter 5). Nonetheless, it bears noting that any *F* could be replaced with an *x* by creating composite variables from the indicators. Although one would gain a slight advantage over multiple regression analysis because the two equations could be estimated simultaneously, one would lose the important benefits of modeling relations between latent rather than observed variables.

Equipped with a basic understanding of the origins of SEM, its relation to statistical models traditionally used by researchers in social and personality psychology, key terminology, and the path diagram, you are now in a position to learn enough about how SEM works and how it can be applied to contemplate using it in your own work. In the remainder of the book, I offer a nontechnical treatment of key features of SEM, presented in the context of a framework for its application. I then review a representative set of models that could be fruitfully applied to the rich conceptual models typical of social and personality psychology.

Note

1 This book was written during a sabbatical leave generously provided by Duke University and funded in part by grant P30 DA023026 from the National Institute on Drug Abuse. I thank Erin Davisson, Cameron Hopkin, and Kaitlin Toner for providing feedback on a draft of the book. The Series Editor, John Nezlek, and Sage London Editor, Michael Carmichael, provided important input on format and style. As with all personal projects, I owe a debt of gratitude to my wife, Lydia, for covering for me so I could focus long enough to finish.