A Mechanism-Based Approach to the Identification of Age-Period-Cohort Models

Citation

Winship, Christopher, and David J. Harding. 2008. A mechanism-based approach to the identification of age-period-cohort models. *Sociological Methods & Research* 36 (3):362-401.

Published Version

<http://dx.doi.org/10.1177/0049124107310635>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3294650>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Sociological Methods & Research

<http://smr.sagepub.com>

A Mechanism-Based Approach to the Identification of Age Period Cohort Models

Christopher Winship and David J. Harding
Sociological Methods Research 2008; 36; 362
DOI: 10.1177/0049124107310635

The online version of this article can be found at:
<http://smr.sagepub.com/cgi/content/abstract/36/3/362>

Published by:

<http://www.sagepublications.com>

Additional services and information for *Sociological Methods & Research* can be found at:

Email Alerts: <http://smr.sagepub.com/cgi/alerts>

Subscriptions: <http://smr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations <http://smr.sagepub.com/cgi/content/refs/36/3/362>

A Mechanism-Based Approach to the Identification of Age-Period-Cohort Models

Christopher Winship

Harvard University, Cambridge, Massachusetts

David J. Harding

University of Michigan, Ann Arbor

This article offers a new approach to the identification of age-period-cohort (APC) models that builds on Pearl's work on nonparametric causal models, in particular his front-door criterion for the identification of causal effects. The goal is to specify the mechanisms through which the age, period, and cohort variables affect the outcome and in doing so identify the model. This approach allows for a broader set of identification strategies than has typically been considered in the literature and, in many circumstances, goodness of fit tests are possible. The authors illustrate the utility of the approach by developing an APC model for political alienation.

Keywords: *APC models; mechanisms; cohorts; front-door criterion; identification strategies*

Age-period-cohort (APC) models are one of the key workhorses used by social scientists in the quantitative analysis of social change. An APC model attempts to decompose temporal change in a dependent variable into period effects, cohort effects, and aging effects, typically using repeated cross-section survey data. A large body of literature going back to the 1970s has examined the problem of identification in APC models (e.g., K. O. Mason et al. 1973; Fienberg and Mason 1979; Glenn 1981; Rodgers 1982; W. M. Mason and Fienberg 1985a). As is well known, without further identifying restrictions, linear and additive APC models are not identified since age (years since birth), period (current year), and cohort (year of birth) are exact linear functions of each other because of the identity $\text{Age} = \text{Period} - \text{Cohort}$.¹

The past literature on the identification of APC models has a number of problems. Beyond the insight that parameter restrictions are needed for identification, the literature has yet to provide a framework for thinking about how APC models might be identified.² Particular parameter restrictions often have not been theoretically well motivated. The results obtained from models also often can be quite sensitive to which parameter restrictions are made (Glenn 1976; Rodgers 1982). Finally, identifying restrictions are rarely, if ever, tested, with the consequence that one must assume that those restrictions are correct.

In this article, we propose a different approach to APC models. Rather than seeing the problem of identification as one of choosing a set of parameter restrictions that are adequate for identification, we frame the problem as one of *theoretically* specifying a model in a sufficiently rich way that it is identified, or better still, overidentified. We propose doing this by specifying the mechanisms by which aging, period-related changes, and cohort-related processes act on the dependent variable. By adding these variables to the model, identification often is possible. In general, it is necessary to fully specify the mechanisms with only one of the APC variables.

Key to our approach is abandoning the goal in much of the previous literature of attempting to find a general, omnibus, *mechanical* procedure for identifying any APC model. Our belief is that this goal is both unattainable and misguided. As Heckman and Robb have stated,

The age-period-cohort effect identification problem arises because analysts want something for nothing: a general statistical decomposition of data without specific subject matter motivation underlying the decomposition. In a sense it is a blessing for social science that a purely statistical approach to the problem is bound to fail. (1985:144-45)

Authors' Note: A version of this article was previously presented at the annual meeting of the American Sociological Association, August 15-19, 2003, Atlanta, Georgia, under the title "The Analysis of Over-Determined Outcomes: Model Identification in the Presence of Functional Dependence." Harding acknowledges support from a National Science Foundation Graduate Research Fellowship and from the Inequality and Social Policy Program at Harvard University, which is funded by a National Science Foundation Integrative Graduate Education and Research Traineeship (IGERT) grant. We are grateful to Glenn Firebaugh, Gary King, Kenneth Bollen, Kenneth Land, Robert Mare, Stephen Morgan, Michael Sobel, and members of Harvard's Applied Statistics Colloquium for comments on earlier drafts of this article. We also thank Cheri Minton for her assistance with the implementation of the bootstrap confidence intervals. Please address correspondence to Christopher Winship, Harvard University, 620 William James Hall, 33 Kirkland Street, Cambridge, MA 02138; e-mail: cwinship@wjh.harvard.edu.

We suggest that what is needed instead is a flexible framework for thinking about the relationship between the particular theoretical model that a researcher has posited and the formal, mathematical conditions that are needed for identification. We offer such an approach. The core idea is that identification can be achieved by extending models to include variables that specify the mechanisms through which age, period, and cohort affect the outcome. As explained in more detail below, the addition of new mechanism variables amounts to expanding an APC model into multiple APC models with the aim of identifying each model and, in doing so, identifying the parameters of interest in the original model.

There is a strong parallel between the logic of our approach and that of instrumental variables (IVs). Understanding this parallel is critical to understanding what we have accomplished and its limitations. In both cases, we can understand the inability to estimate parameters of interest as a problem of model underidentification. Whereas IVs involve adding variables that extend one's model backward to achieve identification, we show how APC models can potentially be identified by adding variables that extend one's model forward, that is, by specifying the different mechanisms through which age, period, and cohort affect the outcome of interest. As in IVs, whether a particular model is identified depends on the theoretical richness of the specification and the availability of measures of specific variables. Thus, as with IVs, in some cases our approach will work and in other cases it will not. However, in contrast to IVs, in many situations the theoretical assumptions underlying the identification strategy of our approach will be testable. Below we discuss in detail the formal identification conditions associated with our approach.

Our approach formally relies on Pearl's (1999, 2000) recent and seminal work on the identification of causal models. Specifically, we show that his front-door criterion provides the basis for identifying separate effects for independent variables that are linearly (or, more generally, functionally) dependent. We demonstrate how his approach can be used to develop a framework for identifying APC models. This provides a number of different strategies for identifying APC models not previously recognized in the literature. Furthermore, we show that in many circumstances model goodness-of-fit tests are available.

Most if not all of the previous methods for the identification of APC models can be formulated within our approach. As in previous work, our approach to identification involves imposing parameter restrictions, though often the restrictions involved may be implicit and may be considerably more complicated than those previously considered. This equivalence is

due to the mathematical requirement that restrictions are needed to identify any APC model (as is true with IVs). This commonality between our approach and previous work, however, should not lead the reader to believe that there are only minor differences between our approach and that in previous work. Specifically, our approach differs because it focuses on the particular theoretical model and the mechanisms that potentially connect age, period, and cohort to the outcome rather than to parameter restrictions. This leads to a distinctly different way of thinking about identification.

Our argument for the importance of mechanisms in APC models complements recent theoretical work in sociology that has argued that sociologists need to pay considerably more attention to specifying the mechanisms through which social processes work (e.g., Hedstrom and Swedberg 1998; Reskin 2003). This work has argued that much sociological theory is too abstract, and to generate testable hypotheses about particular processes, it is necessary to specify the mechanisms involved. For example, Reskin (2003) argues that to test for and understand discrimination, one needs to identify the mechanisms by which it occurs. One cannot simply refer to gender or race differences. This article makes a parallel argument: To achieve identification of APC models, it is necessary to specify the mechanisms through which the processes of interest work.³ The essential point in both the theoretical literature and this discussion is the same: To know why two events are associated, one needs to be able to identify the mechanisms involved.

While we focus on APC models, our approach can be applied to other problems in which there are substantively distinct but linearly or, more generally, functionally dependent explanatory variables. Typically, these are models in which a researcher is interested in two or more variables representing main effects as well as the effect of the difference or sum of these variables. One class of models of this type are "multiple clock" problems such as the APC model. Two other examples of this type are the linear dependence of age, years of work experience, and years of education and the linear dependence of age, age at marriage, and marital duration. This form of linear dependence is also found in other classes of models such as status inconsistency models that attempt to assess the effect of two different statuses and their degree of consistency on some outcome or any of a variety of mobility models that seek to determine the importance of an individual's early and later status or a father's and son's status and the mobility represented by their difference.⁴ More generally, our approach provides a potential solution to any situation where a researcher's *X*s are linearly dependent or nearly so.

In the next section of the article, we discuss previous research. Following this, we briefly discuss Pearl's three criteria for identifying causal effects. We then discuss how APC models can be identified using a mechanism-based approach that draws on these criteria, and we examine different types of APC models. Next, we show how Pearl's front-door criterion can be used to deal with unobserved variables. The subsequent section presents a key goodness-of-fit test. We then present our empirical example. After that, we compare estimates from our model to those using more traditional methods.

Limitations of Previous Research

Although the relationship among age, period, and cohort can be specified in terms of an exact deterministic mathematical relationship, social scientists in general, and sociologists in particular, often argue that they represent three distinct types of social/psychological processes. For example, changes in a dependent variable with respect to age might represent psychological change with age and/or the changing role positions of individuals as they age (e.g., employment, marriage, parenthood, retirement, widowhood, or the empty nest). Changes with respect to period would represent the effects of the current condition of society—for example, if we were referring to the United States, whether the country was in the middle of a war, whether the president was a Republican or a Democrat, or whether the country was in a period of economic boom or recession. Finally, a cohort effect could represent the effect of being born during a specific period (the most famous example is discussed in *The Children of the Great Depression*; Elder 1974) or specific properties of a cohort, such as its size. The problem is that although it is easy to specify distinct social processes related to the general processes associated with age, period, and cohort, it is not possible to straightforwardly estimate the parameters associated with age, period, and cohort because of their linear dependence.

As noted above, the discussion of identification within the technical APC literature has focused on placing restrictions on parameters to identify a model. This is typically done in three ways. First, identification can be achieved if one assumes that only two of the three APC variables affect the outcome. A large number of studies in fact achieve identification by simply assuming that only two of the three variables in an APC model affect the outcome (e.g., Firebaugh and Davis 1988; Glenn 1994; Meyers and Lee 1998). This is a very strong theoretical assumption that may or may not be justified in particular circumstances.

Second, as suggested in K. O. Mason et al. (1973), some set of parameters may be constrained to be equal. For example, based on some theoretical argument it may be assumed that the parameters associated with two periods should be constrained to be equal. This strategy has been used by K. O. Mason et al., Knoke and Hout (1974), Harding and Jencks (2003), and others. More generally, identification might be achieved by assuming that two age, two period, or two cohort parameters are equal. K. O. Mason et al. show that such constraints generally will identify an APC model. The most sophisticated version of this approach has been developed by Nakamura (1986), who uses a Bayesian approach to specify restrictions (for an application, see Sasaki and Suzuki 1987).

A third approach is to constrain the effect of a variable to be proportional to some other substantive variable. For example, it may be assumed that the effect of cohort is proportional to cohort size (Mason and Fienberg 1985b; Kahn and Mason 1987), or a period effect might be restricted to be proportional to the unemployment rate (Farkas 1977). Heckman and Robb (1985) term this the "proxy" variable approach because age, period, and/or cohort are represented by some other variable. O'Brien (2000) terms it the APC-characteristic model. Typically, the proportionality constraint is justified by asserting that the mechanism through which the variable of interest (age, period, or cohort) affects the outcome is captured by the variable used to constrain that variable's effect. O'Brien (2000) provides an advanced discussion of this strategy. The approach in this article, both theoretically and mathematically, generalizes the proxy variable approach. As in the proxy variable approach, the key to identification is specifying the mechanisms through which the APC variables affect the outcome variable. However, unlike the proxy variable approach, we do not make the assumption that the proxy variable is a nonlinear function of age, period, or cohort. In addition, whereas the proxy variable approach assumes that the effect of any APC variable is mediated through only one mediating variable and that mediating variable is affected by only one APC variable, our approach drops both of these restrictions, allowing for a much more general set of models.

Although imposing restrictions certainly provides a solution to identifying the APC model, there are serious problems. First, it is often difficult to find restrictions that can be theoretically justified. Second, if the restrictions are even mildly misspecified, this can have major consequences for parameter estimates (Glenn 1976; Rodgers 1982). Third, restrictions are rarely if ever tested. Typically, this is because the models considered are

just identified, forcing researchers to assume that they have the correct model.⁵

If more constraints are imposed on a model than are needed to just identify it, it is possible to test the adequacy of one's model specification. For example, if a researcher assumes that only one of age, period, or cohort affects the outcome, then this can be tested using standard methods by including either of the other omitted variables as a predictor. In the case where age, period, and cohort each have a proxy variable, one can test one's specification by entering age, period, or cohort into one's models and using standard tests to determine if the inclusion has an effect. More generally, if more restrictions are imposed on a model than are necessary to identify it, it is possible to test whether these restrictions hold. Unfortunately, this type of model specification test is seldom done in the empirical APC literature. Rather, one is typically asked simply to accept that the assumptions that have been made and their related restrictions are valid.

Pearl's Front-Door Criterion

In his 2000 book, *Causality*, Judea Pearl develops a theory for the identification of causal effects in nonparametric models. Pearl's theory uses Bayesian causal networks. He shows that by representing causal relationships between variables in terms of directed acyclic graphs it is possible to use a set of relatively simple graph theoretic criteria to determine when a particular causal model is identified based on a set of observed conditional associations. Key to his thinking is that causal relations represent autonomous mechanisms by which one variable affects another.

We provide a brief overview of Pearl's thinking for two reasons. First, his front-door criterion for identification provides a formal justification for the models presented in this article. Second, his criteria for identification, particularly the backdoor criterion, will be useful in determining the identification status of models and their subcomponents discussed below. The fact that his criteria are easily understood is an added advantage.⁶

The general problem that Pearl (2000) is concerned with is distinguishing true causation from simple statistical association. In his theory, it is assumed that all causal variables, whether observed or unobserved, and the associated causal relations relevant to an outcome are explicitly represented in the graph. To simplify our figures, we omit the error terms.⁷

Figure 1
Backdoor Criterion

Figure 1 shows a simple example in which X and Y are directly connected and also indirectly connected by a path through Z_1 and Z_2 . Pearl (2000) describes three strategies for identifying a causal effect from a set of observed associations.

Pearl's first principle of identification is what he (2000) calls the backdoor criterion. The backdoor criterion amounts to finding variables that when removed from the graph (which is statistically equivalent to conditioning on these variables) cause all pathways between X and Y , other than the direct (causal) path, to be eliminated.⁸ If at least one of the variables in each backdoor path is observed, then the effect of X on Y can be identified. The effect of X on Y is estimated simply by conditioning on one of the variables in each path. This might be done through regression, matching, stratification, or any other conditioning method.

As an example, in Figure 1 the zero-order association between X and Y does not provide an estimate of the effect of X on Y because their association is in part a function of the pathway connecting X and Y through the Z s. Deleting either Z_1 or Z_2 from this graph, which is statistically equivalent to controlling for them, eliminates this pathway. As a result, the conditional association between X and Y now estimates the causal effect of X on Y . Although the example here is extraordinarily simple, the backdoor criterion can be used to prove identification in more complicated situations.

Pearl's second method of identification is the standard IV approach (Pearl 2000). As in Figure 1, the issue is that there are one or more indirect paths connecting X and Y , with the result that the association between X and Y cannot be used to estimate the causal effect of X on Y . The solution with IVs is to augment the model by adding one or more variables that (a) either directly or indirectly affect X , and (b) do not affect Y through

Figure 2
Instrumental Variables

any other pathway. Figure 2 illustrates. The IV Z can be used to identify the effect of X on Y by first estimating the effect of Z on X , the association between Z and Y , and then solving out for the effect of X on Y .

Pearl's third method of identification, the front-door criterion, is likely to be the least familiar to social scientists generally and to sociologists in particular (2000).⁹ The front-door criterion identifies the causal effect of a variable on an outcome by augmenting the causal model to include *all* the intermediate variables through which that variable affects that outcome. If it is possible to identify the effect of the variable of interest on each of the intermediate variables and to identify the effect of each of these variables on the outcome, then the (total) effect of the variable of interest on the outcome can be estimated as the sum of the effects of the paths connecting them.

Pearl uses the example of the effect of smoking on cancer (2000). In Figure 3, we would like to estimate the total effect of S (smoking) on C (cancer). The covariance/correlation between S and C does not provide a consistent estimate because of the backdoor path through U , where U represents possible genetic or environmental factors. If U is observed, then the backdoor criterion shows that we can estimate the effect of S on C by conditioning on U . If U is unobserved, which we represent by enclosing it in an oval, this strategy is not available. However, if we can consistently estimate the effect of S on T (tar) and the effect of T on C , getting estimates of b and c , then we can estimate the effect of S on C as bc . This is the core idea behind the front-door criterion.

In the present case, we can estimate both b and c by a double application of the backdoor criterion (Pearl 2000). Because there are no backdoor paths between S and T , we can consistently estimate the effect of S on T . There is, though, a backdoor path between T and C through S and U .

Figure 3
Front-Door Criterion

However, by conditioning on S , we can eliminate this backdoor path, which allows us to consistently estimate the effect of T on C .

Note that this example assumes that S affects C through T . If there are other mediating variables that are measured, the front-door method can be used to estimate the effects of S on C through these variables. If there are unmeasured mediators, then it will be possible to identify only that component of the effect of S on C that flows through the observed variables. As such, the model is only partially identified.

Identifying APC Models Using a Mechanism-Based Approach

We, as well as others, have thought that Pearl's front-door criterion was an interesting idea but that it would have little application to sociology since it would be too hard to find the intermediary variables, the T s. We argue here, however, that the front-door criterion provides a framework for thinking about the estimation of causal effects when there is linear or functional dependence among our independent variables, exactly the situation in APC models.

The basic idea behind the front-door criterion is to achieve identification by adding variables to one's model that are intermediate between the

independent variables and the outcome variable. By adding variables, we are in an important sense adding additional data to the analysis. These variables would represent the mechanisms through which the original independent variables affect the outcome. The hope is that although the original model is not identified, the subcomponents of the new model will be identified, leading to the full or partial identification of the original model.

Because the augmented model contains intermediate variables, there are now additional endogenous variables besides the final outcome of interest. Associated with each new endogenous variable is an equation with the endogenous variable being a function of age, period, and/or cohort. Thus, each equation is its own APC model. For the overall model to be identified, each equation must be separately identified. Standard identification conditions for APC models apply to each equation in the model. The power of the mechanism-based approach is that the separate mechanism equations will often be easier to identify than the original single APC equation, in which the outcome is a function of all three of A, P, and C.

We now formalize our approach. Define variables and parameters as follows:

$\mathbf{Y}$ is an $n \times 1$ vector measuring the outcome of interest.

$\mathbf{X}$ is an $n \times 4$ matrix consisting of the following variables:

Constant = 1,

Age (A) = years since birth,

Period (P) = current year, and

Cohort (C) = year of birth.

$\mathbf{e}$ is an $n \times 1$ vector for the error term in the APC regression equation.

$\mathbf{a}$ is a 4×1 vector of parameters to be estimated, corresponding to the constant and age, period, and cohort variables.

Our goal is to estimate

$$\mathbf{Y} = \mathbf{X}\mathbf{a} + \mathbf{e} = \mathbf{a}_0 + \mathbf{A}\mathbf{a}_1 + \mathbf{P}\mathbf{a}_2 + \mathbf{C}\mathbf{a}_3 + \mathbf{e}. \quad (1)$$

But because of the identity Age = Period – Cohort, $(\mathbf{X}'\mathbf{X})^{-1}$ does not exist and as a result equation (1) cannot be estimated by ordinary least squares (OLS). The same issue would exist if we allowed age, period, and cohort to have arbitrary nonlinear relationships with $\mathbf{Y}$ by specifying them in

terms of sets of dummy variables. This is what we do in the empirical analysis presented below.

Now generalize the model by assuming that there is a matrix of m variables, $\mathbf{M}$, that is $n \times m$, which represents the mechanisms through which age, period, and cohort affect the outcome $\mathbf{Y}$. Also let $\mathbf{B}$ be a $4 \times m$ matrix of parameters relating a constant, age, period, and cohort to the mechanism variables. Let $\mathbf{U}$ be an $n \times m$ matrix of errors in the equation specifying the relationships between $\mathbf{M}$ and $\mathbf{X}$. Let $\mathbf{c}$ be an $m \times 1$ vector of parameters to be estimated that represent the effect of each mechanism on the outcome $\mathbf{Y}$. We then have the following set of equations:

$$\mathbf{Y} = \mathbf{X}\mathbf{a} + \mathbf{e} \quad \text{Equation specifying the relationship between the Outcome } (\mathbf{Y}) \text{ and Age, Period, and Cohort } (\mathbf{X}). \quad (2)$$

$$\mathbf{M} = \mathbf{X}\mathbf{B} + \mathbf{U} \quad m \text{ equations specifying the relationship among the mechanisms } (\mathbf{M}) \text{ and Age, Period, and Cohort } (\mathbf{X}). \quad (3)$$

$$\mathbf{Y} = \mathbf{M}\mathbf{c} + \mathbf{v} \quad \text{Equation specifying the relationship between the outcome } (\mathbf{Y}) \text{ and the mechanisms } (\mathbf{M}). \quad (4)$$

Substituting (3) into (4), we get

$$\mathbf{Y} = \mathbf{X}\mathbf{B}\mathbf{c} + \mathbf{U}\mathbf{c} + \mathbf{v} \quad \text{Reduced form equation specifying the relationship between } \mathbf{Y} \text{ and Age, Period, and Cohort } (\mathbf{X}). \quad (5)$$

If we can estimate $\mathbf{B}$ and $\mathbf{c}$, then we can estimate $\mathbf{a} = \mathbf{B}\mathbf{c}$, the effects of the APC variables on $\mathbf{Y}$.

Equation (4) above is a linear equation to which standard identification criteria apply. In this case, the $\mathbf{M}$ variables must be linearly independent of each other. Note that if the APC variables are not linearly dependent on the variables in $\mathbf{M}$, then it would be possible to include two of the three APC variables in $\mathbf{M}$, in which case equation (4) would still be identified. **This implies that to achieve full identification, it is necessary to specify the full set of mechanisms associated with only one of the three APC variables.** None of the mechanisms related to the other two APC variables need to be specified, though they could be. Partial identification can be achieved under even weaker conditions. We discuss these issues in more detail below.

The equations in (3) consist of m new APC models. For these equations to be identified, it is sufficient that at least one of A, P, or C be omitted from the model.¹⁰ We consider assumptions of this type below.

Now consider how the above relates to Pearl's front-door criterion. Let M_A be a variable representing the mechanism associated with age, and let

Figure 4
Hypothetical Age–Period–Cohort Model
With Intervening Mechanisms

M_P and M_C be analogously defined. Then we could represent these relationships in terms of the diagram in Figure 4. In the model in Figure 4, we should be able to estimate **b** coefficients since there is no linear dependence problem or, if the relationship is deterministic, we should be able to specify these coefficients. A necessary condition for estimating the **c** coefficients, that is, the effects of M_A , M_P , and M_C , is that the M s be linearly independent. Once we have an estimate of the **b** and **c** coefficients, we can then calculate the relative contribution of age, period, and cohort to the change in the outcome Y as their products. One way to think about the older proxy variable approach is that it is a particular application of Pearl's front-door criterion.

Alternative Types of APC Models

The front-door approach suggests that we can identify the effects of variables by introducing intermediary variables that specify the mechanism(s) by which our variables of interest affect the outcome. In an important sense, it is an extension of the proxy variable approach. The proxy variable approach assumes that there is one distinct and separate variable associated with either age, period, or cohort. That assumption is dropped here, and we allow for the possibility that an intermediary variable may be affected by

Figure 5
Hypothetical Age-Period-Cohort Model
With Shared Mechanisms

any pair of variables consisting of age, period, and cohort. In addition, age, period, and cohort may affect the outcome through several intermediary variables as opposed to only one (as in the standard proxy variable approach). As such, there is a much richer set of models that are identified than those that have typically been considered.

Consider Figure 4 again. This model contains multiple restrictions. First, it assumes that none of the APC variables directly affect the outcome. This amounts to three restrictions. Second, each A , P , and C variable is assumed to affect only one M variable. This amounts to six additional restrictions. Thus, the model in Figure 4 has a total of nine restrictions. As pointed out above, only one restriction is needed to identify an APC model. As a result, more general models that do not contain these restrictions can be considered and, because the model is overidentified, its goodness of fit can be tested.

Figure 5 illustrates the idea that mechanisms may be shared. This model is fully estimable. The effects of A and P on T and similarly the effects of P and C on S can be estimated since they are not linearly dependent on each other. Via the backdoor criterion, the effect of T on Y can be

estimated by conditioning on S , and similarly the effect of S on Y can be estimated by conditioning on T .

There are two basic differences between this model and the standard APC model with proxy variables. First, both T and S are each functions of two variables, not one. The assumption here is that T is affected by age and period and S by period and cohort. Second, period affects both T and S . Because of these two differences, it is difficult, if not impossible, to think about identification as coming from restrictions of the type that have previously been considered in the APC literature. Below, we provide a substantive example in which effects of this type occur.

Identification in the Presence of Unobserved Mechanisms

The problem with consistently estimating any causal effect is the possibility that there are unobserved variables that are associated with both the causal variables and the outcome. We discussed this briefly with regard to Figure 4. In terms of our approach, the concern is that we have not identified all the mechanisms through which age, period, and/or cohort affect the outcome. In this case, we will fail to estimate the total effect of one or more of these variables on the outcome. As described up to this point, the front-door criterion makes the very strong assumption that we have identified all the mechanisms through which age, period, and cohort work. However, in APC models this condition can be relaxed. All that is necessary to identify a model is that we have identified all the mechanisms for *one* of the three APC variables. When this is the case, the effects of the other two APC variables can be controlled for by simply including them directly in the equation predicting the outcome. We illustrate this below.

We now consider the problem of unspecified mechanisms more explicitly. Doing so demonstrates both the power and limitations of Pearl's identification theory, particularly the front-door criterion. Consider Figure 6, which is identical to Figure 4 except that there is an additional path connecting A and Y through an unobserved mechanism variable UM_A and an additional path connecting C and Y through an unobserved mechanism variable UM_C . UM_A and UM_C should be thought of as unspecified or unobserved mechanisms. As before, we enclose these variables in ovals to indicate that they are unobserved. The question is whether we can consistently estimate the total effects of age, period, and cohort on Y or, less ambitiously, whether we can consistently estimate the **b** coefficients. Pearl's front-door criterion states that if we can consistently estimate the **b** and **c** coefficients, then we can consistently estimate the total effects of

Figure 6
Hypothetical Age-Period-Cohort Model
With Unobserved Mechanisms

age, period, and cohort by their appropriate products. For the moment assume that the estimation of the $\mathbf{b}$ coefficients is unproblematic. Also assume that the $\mathbf{M}$ variables are *not* deterministic functions of each other.

For the $\mathbf{c}$ coefficients to be identified, two conditions must hold. First, in whatever conditioning we do, the variable of interest and the conditioning variables cannot be deterministic functions of each other. This is just a more general way of stating the linear dependence problem. Second, we need to be able to break the backdoor paths through UM_A and UM_C connecting each $\mathbf{M}$ variable and Y .

Consider the problem of estimating the effect of M_P on Y , c_2 . There are a variety of backdoor paths between M_P and Y . If there were no unobserved UM variables, as in Figure 4, then c_2 could be consistently estimated by simply conditioning on M_A and M_C by, for example, using a

regression model (as long as they are not deterministic functions of each other). Above, we discussed the conditions necessary for this to be true.

In contrast, in Figure 6, conditioning on M_A and M_C still leaves the backdoor paths $M_P - P - A - UM_A - Y$ and $M_P - P - C - UM_C - Y$. These paths, however, could be eliminated by conditioning on A and C . Since by assumption A , C , and M_P are not exact functions of each other (which would be the case in most empirical applications), the effect of M_P on Y , c_2 , is identified.

Now consider the problem of estimating c_1 . If M_A is a deterministic function of A , then it will not be possible to estimate c_1 conditioning on A . There will be a dependence problem. Let us say, however, that there is variation in M_A independent of A . This would be true for a variable such as education. Other examples would be employment, number of children, or church attendance. Because there is independent variation in M_A , it will be possible to estimate c_1 by conditioning on A . Note that there is no need to condition on either M_P or M_C . Conditioning on A breaks all backdoor paths between M_A and Y .

Assume, however, that the model is a bit more complicated and that M_A is also affected by C . In this case there would now be the backdoor path $M_A - C - M_C - Y$ between M_A and Y . Here, we would need to condition on C as well as A to break all backdoor paths between M_A and Y . In most circumstances M_A , A , and C will not be linearly dependent, and as a result, c_1 will be identified.

The education example shows that there is an additional identification strategy in APC models. Above, we noted that the variable parameterization method, Heckman and Robb's (1985) proxy variable approach, and O'Brien's (2000) APC-characteristic model achieve nonparametric identification by equating the effects of some set of dummy variables or achieve parametric identification by assuming some particular functional relationship between the proxy variable and the outcome. The education example demonstrates that when an intermediary variable contains some variation *independent* of the variables on which it depends, then its effect can also be identified by conditioning on those variables.

Finally, consider the problem of estimating the total effect of A on Y . This is equal to $(b_0c_0) + (b_1c_1)$. Logically, there is no reason that we cannot simply drop M_A and UM_A from the graph in Figure 6 and draw a single line between A and Y that would be equal to this total effect. The question now is whether it is possible to estimate this total effect. There are backdoors between A and Y through both P and C . Conditioning on both P and C is not possible because of the perfect dependence between these three

variables. We could, however, break these backdoor paths by conditioning on M_P and C . It is important to note that this demonstrates that identifying the total causal effects of age, period, or cohort requires specifying the complete set of mechanisms associated with only one of these three variables.¹¹ Thus, in this example, it is possible to identify all three effects if all the effects associated with P are observed.

The requirement that one specify the complete set of mechanisms associated with only one of the APC variables may not be possible in some situations. In this case all may not be lost. Although it may not be possible to fully identify all the coefficients in one's model, it may be possible to identify a subset of the coefficients. Consider Figure 6 again. Assume that P affects UM_A ; that is, there is an arrow going from P to UM_A . In this case there would be unobserved mechanisms associated with all three of the APC variables. Note, however, that it would be still possible to identify the effects of A on Y through M_A , of P through M_P , and of C through M_C by first estimating the effects of M_A , M_P , and M_C on Y and then estimating the effects of the APC variables on each of their respective observed mechanisms. The potential usefulness of these estimates will generally depend on their size and the assumptions one is willing to make about the importance of the pathways that are not estimable. More generally, although a model may not be fully identified, it may be possible to estimate the coefficients with many of the pathways within it. In this case, it may be possible to bound the effects of each of the APC variables even if it is not possible to obtain an actual estimate of each variable's effect. We illustrate this below with respect to our empirical example.

Model Goodness-of-Fit Tests

Since APC models are built on standard statistical models—regression, logit, probit, or Poisson—standard statistical significance tests such as t tests, F tests, log-likelihood ratio tests, and so forth are available to assess the statistical significance of different parameters either individually or as a group. In addition, when APC models are overidentified, it is possible to carry out tests of models of goodness of fit as is typically done in structural equation modeling (Bollen 1989; Bollen and Long 1993) or log-linear or grouped logit analysis (Agresti 1990).

The essential idea in most goodness-of-fit tests is the comparison of a candidate model with a model that fully explains the data of interest, a so-called saturated model. In log-linear analysis, the saturated model would consist of one parameter per cell, perfectly predicting the observed cell

frequencies. In structural equation modeling, the saturated model would perfectly predict all variances and covariances between variables of interest.

In the context of APC models, our interest is in determining whether all the variation in the dependent variable associated with age, period, and cohort are captured by the mechanism variables in the candidate model. Since most empirical models in the literature are just identified, this is true by construction and thus not testable. However, in the mechanism-based approach presented here, the models will often be overidentified, making it possible to test the model's overall fit and the fit of its subcomponents. The importance of goodness-of-fit tests is that they allow researchers to test whether they have specified all the mechanisms involved in the effects of age, period, and cohort on the dependent variable.

An overall goodness-of-fit test is most easily understood by considering its constituent parts. As noted above, the mechanism model consists of a set of APC equations—one equation for the outcome and one equation for each mechanism. Testing the overall fit of the model is equivalent to simultaneously testing the fit of these equations. Recognizing this also reveals what can and cannot be tested. For the fit of a particular equation to be tested, it must be overidentified either by restricting two of age, period, and cohort to have zero effects; by the use of proxy variables; or by multiple parameter restrictions. As is generally the case, testing the overall fit of a model amounts to testing the goodness of fit of only those components that are overidentified.

Testing the fit of any one particular APC equation is straightforward. When observations are categorized into h_A age categories, h_P periods, and h_C cohorts (and when age, period, and cohort categories are constructed so as to be linearly dependent), there are $h = h_A + h_P + h_C - 3$ degrees of freedom associated with the APC variables; thus, the saturated model must contain $h - 1$ APC dummy variables plus a constant to account for all the variation in the outcome associated with age, period, and cohort. This can be accomplished, for example, by specifying a saturated model that consists of a constant, $h_A - 1$ age dummies, $h_P - 1$ period dummies, and $h_C - 2$ cohort dummies (which of age, period, or cohort has two omitted dummy variables is arbitrary).¹² A goodness-of-fit test compares the candidate model to the saturated model. In a linear regression model, this could be done with a standard F test. For logit, probit, or other general linear models, this can be done using a chi-square log-likelihood ratio test. In general, it is best to test the fit of all equations simultaneously. Doing so avoids the multiple testing problem—if one carries out enough tests, by

chance alone some will be significant. If the overall model fails to fit, one can analyze the fit of separate equations to try to identify the source of misfit.

In one special case, the above strategy needs to be modified. As discussed above with respect to education, in some cases there are mechanism variables that vary, in part, independently of age, period, and cohort. This independence provides an important potential source of identification. However, because these variables partially vary independently of age, period, and cohort, the goodness-of-fit test described above needs to be slightly modified. Specifically, in predicting an outcome, such variables need to be included both in the candidate model and in the saturated model.

As an example of formulating goodness-of-fit tests, consider testing the model in Figure 4. The saturated model with which we will compare the overall candidate model represented in Figure 4 includes (in addition to the effects in the candidate model) direct effects from A , P , and C to Y (entered as sets of dummy variables, with one of A , P , or C having two omitted dummies), direct effects from P and C to M_A (with either P or C having two omitted dummies), direct effects from A and C to M_P (with either A or C having two omitted dummies), and direct effects from A and P to M_C (with either A or P having two omitted dummies). Should the overall model fail the goodness-of-fit test, tests of the four individual equations for Y , M_A , M_P , and M_C can be used to discover which parts of the model are causing the failure.¹³ This information can then be used as a guide in further augmenting the model in Figure 4 by adding more mechanism variables, either between A , P , or C and Y or between A , P , or C and the existing mechanism variables.

The advantage of goodness-of-fit tests is that they allow researchers to determine whether the model they have proposed or, more precisely, its overidentified components adequately fit the data. Thus, such goodness-of-fit tests provide a means of testing the assumptions about parameter restrictions associated with each equation when it is overidentified. If the goal is to formulate a model for which all the assumptions about mechanisms can be tested, then the ideal model will include direct effects from at most two of A , P , or C to each of the mechanisms and to Y . This will allow each individual equation in the model to be tested against a saturated model that also additionally includes at least one of A , P , or C , without creating a saturated model that cannot be estimated due to linear dependence problems.

A difficulty with goodness-of-fit tests is that often, particularly with large samples, small deviations between the data and the expected data

under the model can lead to rejection of the model. Thus, it is common for F tests and chi-square log-likelihood ratio tests to reject models that actually fit the data quite well. Many different measures of goodness of fit have been proposed in the structural equations modeling literature. In the example below, we report on two in addition to the standard chi-square log-likelihood ratio test, the sample size adjusted Bayesian information criteria (BIC) and Akaike (1973) information criteria (AIC):

$$\text{Sample Size Adjusted BIC} = -2\log \text{likelihood} + k \ln((n-2)/24),$$

$$\text{Sample Size Adjusted AIC} = -2\log \text{likelihood} + 2k + (2k(k+1))/(n-k-1),$$

where n is the sample size and k is the degrees of freedom. Models with lower values of the BIC and AIC provide a better fit to the data. The difference between the adjusted BIC of the constrained model and the adjusted BIC of the saturated model is one statistic. Negative values favor the constrained model, and positive values favor the saturated model. As a rule of thumb, Raftery (1995) suggests that a BIC difference of less than 2 is weak evidence, a difference of 2-6 is positive evidence, a difference of 6-10 is strong evidence, and a difference of 10 or more is very strong evidence in favor of the model with the lower BIC.

Evidence ratios are based on the adjusted AIC. The evidence ratio is a function of each model's delta, the difference between its AIC and the AIC of the saturated model. Using the delta value for both the comparison model and the saturated model, we can calculate the evidence ratio, which is defined as

$$\text{Evidence ratio} = \exp(-0.5\Delta_{\text{comparison}}) / \exp(-0.5\Delta_{\text{saturated}}).$$

We do not review the theoretical basis for the evidence ratio here, but it can be thought of as the ratio of the evidence in favor of the comparison model over the saturated model. It varies from zero to infinity. The larger the evidence ratio, the better the fit of the comparison model compared to the saturated model. A value of one is produced when the AICs of the two models are the same and neither model is preferred to the other. Burnham and Anderson (1998) provide a description of model testing using BIC and AIC.

We use all three measures for our tests because each has advantages and disadvantages. The chi-square test provides a true statistical test that can be used to determine a level of statistical significance, but it has a small penalty for adding additional parameters and therefore tends to favor less parsimonious models, especially in large samples. The AIC and BIC are not statistical tests but rather provide rough guidelines based on a rule of thumb. However, they allow for sample size adjustments and penalties

Table 1
Variable Descriptions for Political Alienation Example

Name	Description	Source
Political alienation (PA)	1 if agree with statement in text	NES (VCF0609)
Age (A)	29 to 56, in 4-year age groups	NES (VCF0101)
Period (P)	Years: 1956, 1960, 1964, 1968, 1976, 1980	NES (VCF0004)
Cohort (C)	1900-1951, in 4-year birth-year intervals	Kahn and Mason (1987)
Relative cohort size	Percentage of U.S. White males in cohort in year	Table 1, Kahn and Mason (1987)
Unemployment rate	Unemployment rate for U.S. males age 20 + in November of year	CPS
Republican president	1 in years in which sitting president is Republican (1956, 1960)	
Watergate	1 in years after Watergate scandal became public (1976, 1980)	
Employment	1 if currently employed	NES (VCF0118)
Education	Years of schooling constructed from categories less than high school, high school, some college, and college	NES (VCF0110)
Church attendance	Continuous latent variable based on ordered categories never, seldom, often, and regularly	NES (VCF0130, VCF0131)

Note: CPS = Current Population Survey; NES = National Election Studies.

for adding additional parameters. Therefore, they tend to favor more parsimonious models. The BIC has a larger penalty for adding more parameters than the AIC.

Empirical Example

To illustrate these ideas, we conduct a basic analysis of the effects of age, period, and cohort on political alienation (PA). Following Kahn and Mason (1987), we use data from White males surveyed by the National Election Surveys for presidential election years (Sapiro, Rosenstone, and the National Election Studies 2002). Here, PA is measured by whether the respondent agrees or disagrees with the statement “I don’t think public

Table 2
Values of Mechanism Variables With Deterministic
Relationships With Period and Cohort

Period	Unemployment Rate (Percentage)	Watergate	Republican President
1956	3.0	0	1
1960	4.8	0	1
1964	3.1	0	0
1968	1.8	0	0
1976	5.7	1	0
1980	5.8	1	0

Relative Cohort Size (Percentage of White Males in Cohort in Each Year)

Cohort	Period					
	1956	1960	1964	1968	1976	1980
1900-1903	7.06	6.61				
1904-1907	7.98	7.53	6.92			
1908-1911	8.92	8.34	7.93	7.08		
1912-1915	9.41	9.12	8.47	7.85		
1916-1919	9.87	9.54	9.07	8.41	6.62	
1920-1923	10.43	10.17	9.8	9.12	7.69	6.8
1924-1927	10.22	10.15	9.82	9.31	7.81	7.09
1928-1931		9.79	9.31	8.85	7.63	7.06
1932-1935			8.75	8.33	7.26	6.63
1936-1939				8.64	7.6	7.01
1940-1943					8.85	8.24
1944-1947					10.51	9.67
1948-1951						11.35

Note: See variable descriptions in Table 1.

officials care much what people like me think.” Those who agree are coded as one, and those who disagree are coded as zero. Other variables are described in Table 1. We restrict our analyses to married White males age 29 to 56 surveyed in 1956, 1960, 1964, 1968, 1976, and 1980 who have no missing data on any of our variables, leaving an n of 1,705 cases.¹⁴ In all models, age, period, and cohort are entered as sets of dummy variables to avoid assumptions about their functional form.

Table 2 indicates the relationship between each level of period and cohort and their deterministic (nonstochastic) mechanism variables. Since the relationships between period and unemployment rate, between period

Figure 7
Simple Age–Period–Cohort Model for Political Alienation
With Intervening Mechanisms

and Watergate, between period and Republican president, and between cohort and period and cohort size are deterministic, we do not discuss them further. Note that unemployment rate has a highly nonlinear relationship with period.

Figure 7 shows a simple model of the relationships between age, period, cohort, and PA. Because we have a recursive hierarchical model, the errors, which by assumption are independent of each other and the other variables in the diagram, are omitted. Analogous to Figure 4, the figure specifies a single intervening mechanism for each APC variable. As such, Figure 7 can be thought of as a proxy variable model in the sense discussed above. This model makes a number of assumptions. First, it assumes that the effect of period on PA operates entirely through whether one is employed, the effect of age on PA operates entirely through church attendance, and the effect of cohort on PA operates entirely through education. Second, it assumes that there are no causal relationships between variables in the diagram that are not connected by pathways, that is, that cohort does not affect church attendance, and so on. If these assumptions are correct (and we have avoided other common problems such as misspecification of functional form, measurement error, etc.), we can easily estimate the effects of age, period, and cohort on PA using a structural

Table 3
Goodness-of-Fit Analysis for Model in Figure 7

	χ^2 <i>df</i>	χ^2 Test Statistic	χ^2 <i>p</i> value	Difference in Adjusted BIC	AIC Evidence Ratio
Overall model	65	229.60	<.001	-47.59	2.785E-20
Equation for political alienation	22	134.65	<.001	40.83	3.028E-20
Equation for employment	17	31.10	.019	-41.40	5.669
Equation for church attendance	16	56.71	<.001	-11.52	5.666E-06
Equation for education	10	7.14	.712	-35.50	763.745

Note: BIC = Bayesian information criteria; AIC = Akaike information criteria.

equation model. The effects of period, age, and cohort on PA are estimated as b_1c_1 , b_2c_2 , and b_3c_3 , respectively. They can be calculated using methods described in Bollen (1989), Stolzenberg (1979), Fox (1980, 1985), Winship and Mare (1983, 1984), and Xie (1989).¹⁵

We can test whether the model in Figure 7 is misspecified using the goodness-of-fit test described above. Table 3 provides relevant fit statistics. The overall model fails with respect to both the chi-square log-likelihood ratio test ($p < .001$) and the AIC evidence ratio but passes with respect to the difference in the adjusted BICs, which tends to favor more parsimonious models. These results suggest that one should be very cautious about using this model to estimate age, period, and cohort effects.

The remaining rows provide parallel goodness-of-fit tests for the different subcomponents of the model. All the equations fail by at least one criterion except for the education equation. The education equation passes by all three criteria (an insignificant chi-square, a negative BIC difference, and an AIC evidence ratio greater than one). The PA equation fails by all three criteria. Failure implies that additional variables are needed to predict the outcome of interest to fully capture the effects of age, period, and cohort. For example, the fact that the PA equation fails its goodness-of-fit tests means that employment, church attendance, and education do not fully explain the effects age, period, and cohort on PA. Thus, estimates of the total effects of age, period, and cohort on PA based on this model will be incorrect.

Figure 8 represents a much extended and more realistic model of the relationships between *A*, *P*, and *C*, the intervening variables, and PA. PA is directly affected by Watergate, Republican president, employment,

Figure 8
Full Age-Period-Cohort Model for Political Alienation
With Multiple Mechanisms

Table 4
Goodness-of-Fit Analysis for Model in Figure 8

	χ^2 <i>df</i>	χ^2 Test Statistic	χ^2 <i>p</i> value	Difference in Adjusted BIC	AIC Evidence Ratio
Overall model	62	62.84	.446	-201.55	2.60E + 15
Equation for political alienation	20	21.94	.344	-63.35	12,105.77
Equation for employment	21	25.04	.245	-64.51	6,839.56
Equation for church attendance	11	10.47	.488	-36.44	399.40
Equation for education	10	5.41	.862	-37.24	1,843.00

Note: BIC = Bayesian information criteria; AIC = Akaike information criteria.

education, and church attendance. There are two stages of intervening variables in this model since one variable that directly affects PA, employment, is not directly affected by A, P, or C. Further intervening variables include cohort size, unemployment rate, and church attendance. The model is further complicated by the fact that a variable directly affecting PA, education, is also an intervening variable for another variable, employment.

Table 4 provides the analogous goodness-of-fit tests for the model in Figure 8. As before, the first row shows different goodness-of-fit measures for the overall model. The overall model passes the chi-square test, and both the BIC and AIC criteria prefer the constrained model represented by Figure 8. For completeness, we also show the fit statistics for each of the components of the model in Figure 8. In all cases, these equations comfortably pass the goodness-of-fit test by all three criteria.

Note that there are no tests for a number of mechanisms in Table 4. In four cases—Watergate, Republican president, unemployment rate, and cohort size—variables are deterministic functions of age, period, or cohort, so no test is needed.

Table 5 provides estimates for the different equations represented by our preferred model in Figure 8. Space limitations prevent us from discussing all the individual coefficients. Probit coefficients are displayed for the equation for PA because we conceive of this binary variable as representing an underlying continuous variable. Logit coefficients are displayed for employment since whether one is employed or not can be thought of as truly binary. Church attendance coefficients are from an ordinal probit model since this variable is measured by four ordered categories but can be thought of as measuring an underlying latent continuous variable.¹⁶ Years of education is a continuous variable, so its coefficients are from an OLS model. Focusing on the equation for PA, we see that church attendance, education, Watergate, and a Republican president all have substantial and statistically significant effects on PA. The effect of employment on PA seems substantively large, but its large standard error makes it statistically insignificant.

From the coefficients in Table 5, it is possible to calculate the effect of age, period, and cohort along each path. These estimates are shown in Table 6. Because age, period, and cohort are all measured as sets of dummy variables, there is no general age, period, or cohort effect. Rather, these effects depend on the specific values of age, period, and cohort that we chose to compare. Table 6 shows an example calculation comparing those in the 1936 to 1939 cohort surveyed in 1976 with those in the 1908 to 1911 cohort surveyed in 1960.¹⁷ Since we have specified cohort and period, we have also implicitly specified the ages that we are comparing. Our first group is age 37 to 40 and our second is age 49 to 52.

Several results are noteworthy in Table 6. First, our model suggests that period as opposed to cohort or age is the crucial factor in explaining PA. This result is qualitatively consistent with Kahn and Mason's (1987)

Table 5
Estimates of Equations Model in Figure 8

	Political Alienation		Employment		Church Attendance		Years of Education	
	Probit	SE	Logit	SE	Probit	SE	OLS	SE
Employment	-0.185	0.166						
Church attendance	-0.107	0.035						
Education	-0.134	0.016	0.164	0.060				
Watergate	0.494	0.084						
Republican president	-0.358	0.080						
Cohort size			0.354	0.126			-0.100	(0.068)
Unemployment rate			-0.341	0.105				
Year 1956					Reference category			
Year 1960					0.100	0.088		
Year 1964					-0.078	0.083		
Year 1968					-0.315	0.089		
Year 1976					-0.355	0.087		
Year 1980					-0.391	0.096		
Age 29-32					Reference category			
Age 33-36					0.139	0.102		
Age 37-40					0.194	0.098		
Age 41-44					0.254	0.100		
Age 45-48					-0.037	0.102		
Age 49-52					0.122	0.101		
Age 53-56					0.058	0.104		

(continued)

Table 5 (continued)

	Political Alienation		Employment		Church Attendance		Years of Education	
	Probit	SE	Logit	SE	Probit	SE	OLS	SE
Cohort 1900-03							Reference category	
Cohort 1904-07							0.256	0.378
Cohort 1908-11							1.126	0.374
Cohort 1912-15							0.956	0.372
Cohort 1916-19							1.250	0.380
Cohort 1920-23							1.837	0.380
Cohort 1924-27							1.825	0.367
Cohort 1928-31							1.989	0.365
Cohort 1932-35							2.219	0.365
Cohort 1936-39							2.223	0.384
Cohort 1940-43							2.733	0.407
Cohort 1944-47							2.987	0.461
Cohort 1948-51							2.811	0.568
Constant	1.364	(0.251)	-.372	(1.377)	-1.231	(0.091)	11.788	0.573
Threshold 1					-0.067	(0.087)		
Threshold 2					0.287	(0.087)		
Threshold 3								

Note: $n = 1,705$. OLS = ordinary least squares.

Table 6
Example Calculation of Total Age, Period,
and Cohort Effects Based on Model in Figure 8^a

	Estimate	95 Percent CI
Period effect (1976 vs. 1960)		
P → Watergate → PA	.4939	.3337, .6629
P → Republican president → PA	.3576	.2010, .5205
P → unemployment rate → employment → PA	.0021	-.0017, .0067
P → cohort size → employment → PA	.0018	-.0015, .0059
P → cohort size → education → PA	-.0099	-.0238, .0043
P → cohort size → education → employment → PA	-.0001	-.0003, .0001
P → church attendance → PA	.0487	.0163, .0973
Total	.8940	.7252, 1.0633
Cohort effect (1936-1939 vs. 1908-1911)		
C → cohort size → employment → PA	.0018	-.0015, .0059
C → cohort size → education → PA	-.0099	-.0238, .0043
C → cohort size → education → employment → PA	-.0001	-.0003, .0001
C → education → PA	-.1470	-.2321, -.0676
C → education → employment → PA	-.0012	-.0044, .0009
Total	-.1565	-.2446, -.0791
Age effect (37-40 vs. 49-52)		
A → church attendance → PA	.0077	-.0157, .0310
Total	.0077	-.0157, .0310
Grand total	.7453	.5590, .9214

Note: Metric is continuous latent variable from political alienation probit model, and 95 percent confidence intervals were calculated by bootstrap (1,000 replications).

a. Compare those born from 1936 to 1939 and surveyed in 1976 with those born from 1908 to 1911 and surveyed in 1960.

analysis, which involved a similar but not identical time period and a sample that included a wider age range.

Unlike a traditional analysis, however, the mechanism-based approach provides a direct way of understanding why period affects PA. As can be seen in Table 6, Watergate and the president being Republican are key period-related factors. These two results are hardly surprising. However, Table 6 shows that the period effect also works in part through its effect on church attendance, though this effect is smaller than Republican president and Watergate.

Finally, in contrast to Kahn and Mason's (1987) argument that there are no cohort effects for PA, Table 6 suggests that cohort has a significant negative effect. Kahn and Mason are primarily interested in the effect of

cohort size on PA. In our Table 6, all the effects of cohort that are mediated through cohort size are quite small and statistically insignificant, consistent with their findings. However, the results in Table 6 suggest that Kahn and Mason have missed an important factor—cohort's effect on PA through educational attainment. In particular, in our analysis the cohort–education–PA pathway is negative, moderate in size, and statistically significant, indicating that older cohorts, net of other factors, are more politically alienated because they are less educated. In fact, Kahn and Mason (Table 6, Model 8) present results for a period/cohort model in which age effects are constrained to zero where they also find a negative cohort effect. They, however, reject this model on theoretical grounds, stating, “we are unaware of any reason for supposing that the oldest cohort should be most alienated” (p. 164). Apparently, they did not consider the possible importance of differences in educational attainment as an explanation for cohort differences in PA. Because they assume there is no cohort effect, they underestimate the size of their period effects. This example nicely illustrates how fragile theoretical reasoning can be and how important it is to be able to empirically test theoretical assumptions.

As noted earlier, even when it is not possible to fully identify a model, it may be possible to partially identify it. In the model in Figure 8, we have assumed that period and age affect church attendance, but not cohort. Assume instead that all three variables affect church attendance. In this case, because of the standard linear dependence problem, we would not be able to estimate the pathways going from the three APC variables through church attendance. If, however, we were willing to assume that all three pathways were positive (or negative),¹⁸ then we could bound our estimates. From Table 6 we see that the effect of the APC variables through church attendance is the effect of period through church attendance (.0487) and the effect of age through church attendance (.0077) for a total effect of $.0487 + .0077 = .0564$. Because of the linear dependence problem, we have no sure way of assigning this effect. We can, however, estimate a range for the total effects of age, period, and cohort if we respectively assigned all or none of this effect to each of these variables separately. In this case, the total effect of age would fall in the interval $[0, .0564]$, of period in the interval $[-.8453, .9017]$, and of cohort in the interval $[-.1565, -.1001]$. Note that in this case our qualitative conclusion that period effects are by far the most important factor in explaining changes in PA would remain unchanged. Often, a partially identified model may be quite informative if the bounds on the effects of the APC variables are relatively narrow.

A full analysis of PA using the mechanism-based approach warrants a separate article. The point of the analysis here is to demonstrate how a mechanism-based approach can be carried out and to show the types of insights that are possible with our methodology but that are not available from traditional approaches.

Comparison With Traditional Methods

We have offered a mechanism-based approach as an alternative to traditional methods for the identification and analysis of APC models. Above, we argue that a key advantage of the mechanism-based approach is that many, if not necessarily all, of the assumptions in a mechanism model are potentially testable. This is in sharp contrast to traditional approaches in which models are typically just identified and the reader must assume that the identifying restrictions made by the researcher are correct. In the last section, we also showed how a mechanism-based approach could provide considerably more insight than traditional methods by identifying how age, period, or cohort affect an outcome. Despite these strong advantages, the reader may still want to know whether it makes a difference which method is used for the estimates of the total age, period, and cohort effects. Here we compare the model represented in Figure 8 with various traditional models. We focus on the differences across models in the relative importance of age, period, and cohort for PA.

Table 7 reports estimates for the effects of age, period, and cohort for eight different models. Confidence intervals for these estimates are also reported. The first column summarizes the results from our mechanism-based model reported in Table 6. The second column reports the results from the proxy model represented by Figure 7. Columns 3 through 8 report results from traditional models with a single equation for PA. Columns 3 through 5, respectively, report the results when either cohort, age, or period are assumed to have no effect and are dropped from the model. Columns 6 through 8 report the results from models where period, age, or cohort parameter constraints, respectively, have been imposed. Model 6 restricts two period dummies to be equal, selecting the two consecutive years that have the most similar unadjusted mean levels of PA. Model 7 restricts two age dummies to be equal, selecting the two consecutive age groups that have the most similar unadjusted mean levels of PA. Model 8 restricts two cohort dummies to be equal, selecting the two consecutive cohort groups that have the most similar unadjusted mean levels of PA.

Table 7
Comparison of Mechanism Model Results to Traditional Model Results (Total Effects of Period, Age, and Cohort on Political Alienation for Two Example Age-Period-Cohort Groups)

	Traditional Models							
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Mechanism Model (Figure 8)	Proxy Model (Figure 7)	Period and Age Model	Period and Cohort Model	Age and Cohort Model	Period Constraint Model	Age Constraint Model	Cohort Constraint Model	
Period effect (1976 vs. 1960)	0.894 (0.73, 1.06)	0.008 (-0.00, 0.03)	0.844 (0.61, 1.08)	0.740 (0.47, 1.01)	0	1.062 (-0.05, 2.17)	1.254 (0.15, 2.35)	0.683 (-0.31, 1.68)
Cohort effect (1936-1939 vs 1908-1911)	-0.157 (-0.24, -0.08)	-0.103 (-0.16, -0.05)	0	0.263 (-0.13, 0.66)	1.093 (0.70, 1.49)	-0.258 (-2.04, 1.53)	-0.593 (-2.46, 1.28)	0.406 (-1.33, 2.14)
Age effect (37-40 vs. 49-52)	0.008 (-0.02, 0.03)	-0.012 (-0.05, 0.02)	0.127 (-0.10, 0.35)	0	-0.265 (-0.52, -0.01)	0.332 (-0.47, 1.13)	0.476 (-0.45, 1.40)	0.048 (-0.75, 0.84)

Note: Metric is the continuous latent variable from the political alienation probit model. The 95 percent confidence intervals are in parentheses. Confidence intervals for Models 1 and 2 were calculated by bootstrap (1,000 replications). Effects for the mechanism model were calculated in Table 6.

In Table 7 it is remarkable how sensitive the effects of age, period, and cohort are to the model that is used. Period has a large, positive, and significant effect in Models 1, 3, 4, and 7; a large, positive, but insignificant effect in Models 6 and 8; and a small, negative, and insignificant effect in Model 2. Similarly, the estimates of the cohort effect vary considerably. In Models 1 and 2 it has a moderate, negative, and statistically significant effect. In Model 5 cohort has a large, positive, and significant effect, and in Model 8 it has a large, positive, and insignificant effect. In Models 6 and 7 the effect is negative and insignificant; in Models 4 and 8 it is moderate, positive, but insignificant. Finally, the age effects also vary across models. In Model 5 its effect is moderate, negative, and significant; in Models 1 and 2 it is close to zero and insignificant; and in Models 3, 6, 7, and 8 its effect is positive and insignificant. A multitude of other models could be estimated giving additional results. Almost certainly, there is also some traditional model that gives the same results as our Model 1. However, without Model 1 as a guide, the researcher would be forced to rely solely on theory to choose from among the traditional models since it is not possible to empirically test the assumptions of Models 3 through 8, as we did above for Model 1.

These results show that it matters enormously which APC model one uses in analyzing these data. One can get whatever result one wants by choosing the appropriate model! Because of this, it is critical to have strong reasons for preferring one model to another. Theory can certainly be helpful, but it is better to be able to test the goodness of fit of one's model and thus to test one's assumptions. We strongly prefer Model 1 (Figure 8) to the other models because its identifying assumptions have been tested using our goodness-of-fit tests.

Conclusion

Although there is a large literature on the identification of APC models, to date it has not provided a fully satisfactory solution. In this article we have presented a new methodological approach to the identification of APC models. Our method also is applicable to other multiple-clock models, models of social mobility and status inconsistency, or any model in which there is linear or more general functional dependence or near dependence among variables.

We presented an illustrative example to demonstrate how mechanism-based methods can be used. Specifically, our example illustrates an

explicit theoretical strategy for identifying APC models. This strategy involves specifying the mechanisms by which age, period, and cohort affect the dependent variable. This approach points to a much broader set of identification strategies than has previously been considered. It is possible to have models in which more than one mechanism is associated with age, period, or cohort or models in which age, period, or cohort share a mechanism. It is also possible to have models in which mechanisms contain a component that is independent of age, period, and cohort, providing a previously unrecognized source of identification.

By considering more complicated APC models, it is also possible to test the overall goodness of fit of a model and its subcomponents. Such tests are critical in that they allow researchers to test the plausibility of their assumptions and associated model specification. Finally, mechanism-based APC models reveal the processes underlying age-, period-, and cohort-related changes in social phenomena. As such, they provide deeper insight into why change has occurred.

Like all methods, those proposed here are not without limitations. As with IVs, for example, mechanism-based APC models require appropriate data, in this case a relatively rich set of potential mechanism variables. It also will not always be possible to test all identifying assumptions using goodness-of-fit tests if one or more parts of the model are not overidentified. Nevertheless, we hope that these models will renew interest in and allow more rigorous analysis of APC models and other models in which linear (or functional) dependence is an issue.

Notes

1. Of course, not all age–period–cohort (APC) models are linear and additive. For example, one might specify one of age, period, or cohort as A^2 , P^2 , or C^2 , thus avoiding the identification problem. The approach we develop here, however, allows for the least restrictive functional form possible, specifying all three of age, period, and cohort as sets of dummy variables.

2. In a series of papers, Yang and Land have proposed a new and highly technical solution to the identification of APC models in tables of rates or proportions. See Yang, Fu, and Land (2004), Yang and Land (2006, 2008 [this issue]), and Yang (2006). Whether their approach provides a basis for a general solution to the APC problem has yet to be determined. The relationship between their models and those proposed here is a topic for future research.

3. Methodologists working on APC models have always advocated the use of theory in the identification of such models, but previously theory was used primarily to justify excluding either age, period, or cohort from the model; setting two or more coefficients to be equal; or using a particular proxy variable. Our approach departs from this previous work in *how*

theory is used. Here we argue that theory should be used to identify the mechanisms through which age, period, and cohort have their effects.

4. We are indebted to Robert Mare for this observation.

5. A fourth issue that we do not examine here is whether APC models should be thought of as causal models. Holland (1986) has argued that it makes sense to talk about causality only when a variable is manipulable (for a discussion of this issue and less extreme positions taken by others, see Winship and Sobel 2004). Clearly, an individual's age, the cohort into which he or she was born, or the present period are not manipulable variables. An important literature in philosophy, however, has argued that what is critical to causal analysis is the specification of the mechanism or mechanisms through which a particular causal effect is to occur, not manipulability. This line of reasoning has ancient roots going back to Aristotle's notion of an efficient cause. The key idea is that a cause must have the ability to bring about an effect (Harré 1972; Harré and Madden 1975; Bunge 1979). This is posited to occur because the cause is related to the outcome through some set of mechanisms (Cartwright 1989; Glennan 1996). As the examples we provide later illustrate, the mechanisms through which age, period, and cohort potentially work typically are manipulable, suggesting that APC models can be thought of as causal.

6. In many ways, Pearl's theory is similar to the standard theory of linear path models of Wright (1921) and developed within sociology by Duncan (1975). However, it differs from this theory in three critical respects. First, it deals with nonparametric models of causal effects. Second, it provides a more general theory for the identification of causal effects than that in the standard theory. Third, Pearl explicitly shows the relationship between his theory and the counterfactual model of causal effects (Pearl 1999, 2000).

7. Although our approach is applicable to a wide range of models, in the examples below we consider only recursive hierarchical models. As a result, the errors in our models are assumed to be independent of each other and all the variables, either observed or unobserved, in our models. In this case, nothing is lost by omitting these variables from the diagram.

8. More precisely, all backdoor paths that are not blocked need to be eliminated. A path is blocked if it contains a collider variable. A collider variable is a variable with two or more arrows going into it (Pearl 2000).

9. Sociologists of stratification will recognize the idea of specifying intervening variables to capture causal mechanisms in the Wisconsin model of status attainment (e.g., Sewell, Haller, and Portes 1969; Sewell and Hauser 1980).

10. Identification could also be achieved by imposing parameter restrictions, adopting a proxy variable approach, or by treating the mechanism as the outcome of interest and then specifying the mechanisms through which A , P , and C affect it.

11. Note also that this implies that traditional proxy variable models that use a single proxy variable to capture the effects of age, period, or cohort will fail to correctly estimate APC effects when that single proxy variable does not represent all the causal mechanisms through which that APC variable affects the outcome. Although it is seldom done in the empirical literature, this assumption will often be testable.

12. Note that this saturated model is the model suggested by K. O. Mason et al. (1973) in which identification is achieved by setting two periods, two ages, or two cohorts to be equal. As discussed above, this model has been criticized because different constraints produce the same predicted values but different parameter estimates. However, here we are not interested in the parameter estimates of the saturated model, only its ability to predict the outcome.

13. The saturated model to test against the equation for Y would include the three $\mathbf{M}$ variables and A , P , and C (entered as sets of dummy variables, with one of A , P , or C having two omitted dummies). The saturated model to test against the equation for M_A would include A , P , and C (with either P or C having two omitted dummies). The saturated model to test against the equation for M_P would include P , A , and C (with either A or C having two omitted dummies). The saturated model to test against the equation for M_C would include C , P , and A (with either P or C having two omitted dummies).

14. Kahn and Mason (1987) also include those surveyed in 1952 and 1972, but some variables for our example are not available for those years.

15. Calculating indirect and total effects from structural equation models is most complicated when intervening variables are not normally distributed, as is the case here for employment. Since we treat this variable as a discrete variable (rather than a binary measure of an underlying continuous variable) and estimate its equation using a logit specification, the contribution of an independent variable through employment along an indirect effect path is $Bp(1-p)$, where B is the logit coefficient for the independent variable in the equation predicting employment (see Winship and Mare 1983:85-86). Here we choose p as the mean value of employment in the sample, $p = .962$.

16. When church attendance is entered as an independent variable in the PA equation, we use its predicted values from its ordinal probit equation (constructed from the linear prediction and the estimated thresholds). The correlation between observed and predicted church attendance is greater than .9. This puts church attendance in the metric of the underlying continuous latent variable, greatly simplifying calculation of indirect effects.

17. Table 6 provides estimates of indirect and total effects as well as 95 percent confidence intervals for those estimates. Confidence intervals are calculated by bootstrapping with 1,000 replications. Because they are the product of two or more variables, indirect effects are often not normally distributed, biasing standard errors calculated under an assumption of normality (MacKinnon, Lockwood, and Williams 2004; Shrout and Bolger 2002). Nonnormality seems to be the case here, as many of the confidence intervals are not symmetric around their point estimates.

18. If the paths are of different signs, then the size of any one path will be unbounded. The only constraint here is that the sum of the pathways must be equal to their total effect. In the example, this is .0564.

References

- Agresti, Alan. 1990. *Categorical Data Analysis*. New York: John Wiley.
- Akaike, Hirotugu. 1973. "Information Theory as an Extension of the Maximum Likelihood Principle." Pp. 267-81 in *Second International Symposium on Information Theory*, edited by B. N. Petrov and F. Csaki. Budapest, Hungary: Akademiai Kiado.
- Bollen, Kenneth A. 1989. *Structural Equations With Latent Variables*. Wiley Series in Probability and Mathematical Statistics. New York: John Wiley.
- Bollen, Kenneth A. and J. Scott Long, eds. 1993. *Testing Structural Equation Models*. Newbury Park, CA: Sage.
- Bunge, Mario A. 1979. *Causality and Modern Science*. 3d ed. New York: Dover.

- Burnham, Kenneth P. and David R. Anderson. 1998. *Model Selection and Inference: A Practical Information-Theoretic Approach*. New York: Springer-Verlag.
- Cartwright, Nancy. 1989. *Nature's Capacities and Their Measurement*. New York: Oxford University Press.
- Duncan, Otis D. 1975. *Introduction to Structural Equation Models*. New York: Academic Press.
- Elder, Glen. 1974. *Children of the Great Depression: Social Change in Life Experience*. Chicago: University of Chicago Press.
- Farkas, George. 1977. "Cohort, Age, and Period Effects Upon the Employment of White Females: Evidence for 1957-1968." *Demography* 14:33-42.
- Fienberg, Stephen E. and William M. Mason. 1979. "Identification and Estimation of Age-Period-Cohort Models in the Analysis of Discrete Archival Data." *Sociological Methodology* 10:1-67.
- . 1985. "Specification and Implementation of Age, Period, and Cohort Models." Pp. 45-88 in *Cohort Analysis in Social Research*, edited by W. M. Mason and S. E. Fienberg. New York: Springer-Verlag.
- Firebaugh, Glenn and Kenneth E. Davis. 1988. "Trends in Antiblack Prejudice, 1972-1984: Region and Cohort Effects." *American Journal of Sociology* 94:251-72.
- Fox, John. 1980. "Effect Analysis in Structural Equation Models." *Sociological Methods & Research* 9:3-28.
- . 1985. "Effect Analysis in Structural Equation Models 2: Calculation of Specific Indirect Effects." *Sociological Methods & Research* 14:81-95.
- Glenn, Norval D. 1976. "Cohort Analysts' Futile Quest: Statistical Attempts to Separate Age, Period and Cohort Effects." *American Sociological Review* 41:900-904.
- . 1981. "The Utility and Logic of Cohort Analysis." *Journal of Applied Behavioral Science* 2:247-57.
- . 1994. "Television Watching, Newspaper Reading, and Cohort Differences in Verbal Ability." *Sociology of Education* 67:216-30.
- Glennan, Stuart S. 1996. "Mechanisms and the Nature of Causation." *Erkenntnis* 44:49-71.
- Harding, David J. and Christopher Jencks. 2003. "Changing Attitudes Toward Pre-Marital Sex: Cohort, Period, and Aging Effects." *Public Opinion Quarterly* 67:211-26.
- Harré, Rorn. 1972. *The Philosophies of Science*. Oxford, UK: Oxford University Press.
- Harré, Rorn and Edward H. Madden. 1975. *Causal Powers: A Theory of Natural Necessity*. Oxford, UK: Basil Blackwell.
- Heckman, James and Richard Robb. 1985. "Using Longitudinal Data to Estimate Age, Period, and Cohort Effects in Earnings Equations." Pp. 137-50 in *Cohort Analysis in Social Research*, edited by W. M. Mason and S. E. Fienberg. New York: Springer-Verlag.
- Hedstrom, Peter and Richard Swedberg, eds. 1998. *Social Mechanisms: An Analytical Approach to Social Theory*. Cambridge, UK: Cambridge University Press.
- Holland, Paul W. 1986. "Statistics and Causal Inference (With Comments)." *Journal of the American Statistical Association* 81:945-70.
- Kahn, Joan R. and William M. Mason. 1987. "Political Alienation, Cohort Size, and the East-erlin Hypothesis." *American Sociological Review* 52:155-69.
- Knoke, David and Michael Hout. 1974. "Social and Demographic Factors in American Political Party Affiliations, 1952-1972." *American Sociological Review* 39:700-713.

- MacKinnon, David P., Chondra M. Lockwood, and Jason Williams. 2004. "Confidence Limits for the Indirect Effect: Distribution of the Product and Resampling Methods." *Multivariate Behavioral Research* 39:99-128.
- Mason, Karen O., H. H. Winsborough, William M. Mason, and W. Kenneth Poole. 1973. "Some Methodological Issues in Cohort Analysis of Archival Data." *American Sociological Review* 38:242-58.
- Mason, William M. and Stephen E. Fienberg, eds. 1985a. *Cohort Analysis in Social Research*. New York: Springer-Verlag.
- . 1985b. "Introduction: Beyond the Identification Problem." Pp. 1-8 in *Cohort Analysis in Social Research*, edited by W. M. Mason and S. E. Fienberg. New York: Springer-Verlag.
- Myers, Dowell and Seong Woo Lee. 1998. "Immigrant Trajectories into Home Ownership: A Temporal Analysis of Residential Assimilation." *International Migration Review* 32(3): 593-625.
- Nakamura, T. 1986. "Bayesian Cohort Models for General Cohort Table Analysis." *Annals of the Institute of Statistical Mathematics* 38 (Part B): 353-70.
- O'Brien, Robert M. 2000. "Age Period Cohort Characteristic Models." *Social Science Research* 29:123-39.
- Pearl, Judea. 1999. "Graphs, Structural Models, and Causality." Pp. 95-139 in *Computation, Causation, & Discovery*, edited by C. Glymour and G. F. Cooper. Cambridge, MA: MIT Press.
- . 2000. *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press.
- Raftery, Adrian E. 1995. "Bayesian Model Selection in Social Research." *Sociological Methodology* 25:111-63.
- Reskin, Barbara. 2003. "Including Mechanisms in Our Models of Ascriptive Inequality." *American Sociological Review* 68:1-21.
- Rodgers, Willard L. 1982. "Estimable Functions of Age, Period, and Cohort Effects." *American Sociological Review* 47:774-87.
- Sapiro, Virginia, Steven J. Rosenstone, and the National Election Studies. 2002. *American National Election Studies Cumulative Data File, 1948-2000* (11th ICPSR Version). [Computer file, ICPSR No. 8475]. Ann Arbor, MI: University of Michigan, Center for Political Studies [Producer]. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [Distributor].
- Sasaki, Masamichi and Tatsuzo Suzuki. 1987. "Changes in Religious Commitment in the United States, Holland, and Japan." *American Journal of Sociology* 92:1055-76.
- Sewell, William H., Archibald O. Haller, and Alejandro Portes. 1969. "The Educational and Early Occupational Attainment Process." *American Sociological Review* 34:82-92.
- Sewell, William H. and Robert M. Hauser. 1980. "The Wisconsin Longitudinal Study of Social and Psychological Factors in Aspirations and Achievements." *Research in Sociology of Education and Socialization* 1:59-99.
- Shrout, Patrick E. and Niall Bolger. 2002. "Mediation in Experimental and Non-Experimental Studies: New Procedures and Recommendations." *Psychological Methods* 7:423-45.
- Stolzenberg, Ross M. 1979. "The Measurement and Decomposition of Causal Effects for Nonadditive Models." Pp. 459-99 in *Sociological Methodology 1980*, edited by K. F. Schuessler. San Francisco: Jossey-Bass.
- Winship, Christopher and Robert D. Mare. 1983. "Structural Equations and Path Analysis for Discrete Data." *American Journal of Sociology* 89:54-110.

- . 1984. "Regression Models With Ordinal Variables." *American Sociological Review* 49:512-25.
- Winship, Christopher and Michael Sobel. 2004. "Causal Inference in Sociological Studies." Pp. 481-503 in *Handbook of Data Analysis*, edited by M. Hardy and A. Bryman. Thousand Oaks: Sage.
- Wright, Sewall. 1921. "Correlation and Causation." *Journal of Agricultural Research* 20:557-85.
- Xie, Yu. 1989. "Structural Equation Models for Ordinal Variables: An Analysis of Occupational Destination." *Sociological Methods & Research* 17:325-52.
- Yang, Yang. (2006). "Bayesian Inference for Hierarchical Age-Period-Cohort Models of Repeated Cross-Section Survey Data." *Sociological Methodology* 36:39-74.
- Yang, Yang, Wenjiang Fu, and Kenneth C. Land. 2004. "A Methodological Comparison of Age-Period-Cohort Models: The Intrinsic Estimator and Conventional Generalized Linear Models." *Sociological Methodology* 34:75-110.
- Yang, Yang and Kenneth C. Land. 2006. "A Mixed Models Approach to Age-Period-Cohort Analysis of Repeated Cross-Section Surveys: Trends in Verbal Test Scores." *Sociological Methodology* 36:75-97.

Christopher Winship is the Diker-Tishman Professor of Sociology at Harvard University and also a faculty member in the Kennedy School of Government. He recently published with Steve Morgan *Counterfactuals and Causal Inference: Methods and Principles for Social Research* (Cambridge, 2007).

David J. Harding is an assistant professor of sociology and an assistant research scientist at the Population Studies Center at the University of Michigan. He recently published "Cultural Context, Sexual Behavior, and Romantic Relationships in Disadvantaged Neighborhoods" (*American Sociological Review*, June 2007).