A General Strategy for the Identification of Age, Period, Cohort Models: A Mechanism Based Approach

Christopher Winship

and

David J. Harding

Harvard University

November 2003

A version of this paper was previously presented at the Annual Meetings 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.” The authors can be contacted at Department of Sociology, Harvard University, 33 Kirkland Street, Cambridge, MA 02138 or winship@wjh.harvard.edu and dharding@wjh.harvard.edu. 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.

Abstract

To date, the suggested strategy for identification of the APC model has been to specify some set of restrictions on the effects of Age, Period, and Cohort that allow the model to be identified. Generally, this is done by either restricting a set of coefficients to be equal or by requiring the effect of a variable to be proportional to some other variable. Although imposing restrictions certainly provides a solution to identifying the APC model, existing literature expresses dissatisfaction. First, it is often difficult to find restrictions that can be theoretically justified. Second, if the restrictions are even mildly wrong, this can have a major effect on parameter estimates (Glenn 1976). Third, apart from identification issues, much research has been criticized for being atheoretical and solely interested in decomposing the relative effects of Age, Period, and Cohort. Fourth, the previous literature has not examined what it means to specify an APC model as causal model. An individual's Age, Period, or Cohort cannot be changed and, as such are not manipulable. Holland (1986) and others have argued that in this case one cannot talk about a variable having a causal effect. We suggest that it is meaningful to talk about causal effects in APC models if one can specify the mechanisms that are involved and these mechanisms are manipulable, which typically will be the case..

This paper offers an alternative approach to the problem of identification. Building on Pearl's (1999, 2000) work on nonparametric causal models and in particular his front-door criterion for the identification of causal effects, we provide a formal theory for identifying APC models. The core of the estimation strategy is to specify the mechanisms through which the variables of interest are suppose to work. In contrast to previous research that has focused on identification through restrictions, our method involves identification by augmenting one's model to include the variables that specify the mechanisms by which the variables of interest are assumed to affect the outcome. In a restricted set of cases, the two approaches are mathematically equivalent.

Our approach allows for a much broader set of identification strategies than has typically been considered in the literature. Also, in certain circumstances, different model specification tests are possible. Most importantly, our approach demonstrates the importance of theory for the identification of APC models, particularly theory as the specification of the mechanisms by which different processes are assumed to work. Although our focus is on APC models, our approach is applicable to theory-based evaluation and comparative historical analysis. We do not pursue these investigations here. We illustrate the utility of our approach by developing an APC model for political alienation.

Introduction

Age, Period, Cohort (APC) models are one of the key workhorses used by social scientists in the quantitative analysis of social change. A large literature going back to the 1970's has examined the problem of identification in APC models (e.g. Mason et al. 1973, Glenn 1981, Fienberg and Mason 1979, Mason and Fienberg 1985a). As is well known, without further identifying restrictions, APC are not identified since Age, Period, and Cohort are exact linear functions of each other. For example, Age = Period – Cohort.

The past literature on the identification of APC models has a number of problems and is not entirely satisfactory. Beyond the insight that parameter restrictions are needed, the literature has yet to provide a general framework for thinking about how APC models might be identified. In addition, the focus on particular parameter restrictions has often not been theoretically well motivated. The results obtained from models also can often be quite sensitive to which parameter restrictions are made, and as such sensitive to misspecification of those restrictions (Glenn 1976). The literature has paid little if any attention to testing the validity of particular restrictions. As a result, the restrictions imposed often appear ad hoc and arbitrary. Finally, to the degree that the goal has been to find an omnibus, mechanical procedure for identifying APC models, the search has, at least up until now, failed.

In this paper, 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 identifying an APC model, we frame the problem as one of *theoretically* specifying a model in a sufficiently rich way that it is identified, or better yet overidentified. We argue that one quite general way of doing this is by specifying the mechanisms by which aging, period-related changes, and cohort-related processes act on the dependent variable of interest. Whether a model is identified or not, then, is a function of how rich our specification in terms of mechanisms is.

Key to our approach is abandoning the goal of 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. (Heckman and Robb 1985a)

We suggest that what is needed instead is a powerful 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 identifiability. We offer such an approach. The core idea is that identification can be achieved by extending models to include those variables which specify the mechanisms through which Age, Period, and Cohort are assumed to work.

There is a strong parallel between the logic of our approach and that of instrumental variables. Understanding this parallel is critical to understanding what we have and have not accomplished. In both cases, the issue can be understood as a problem of not being able to estimate the parameters of interest because of model underidentification. A general and universal concern in any regression analysis whose goal is the estimation of one or more causal effects is that one or more independent variables in the regression equation might be correlated with the error term. It is well known that if one or more independent variables are correlated with the error term then ordinary least square estimates (OLS) of all the slope parameters will typically be inconsistent.

One general strategy for dealing with situations where one or more independent variables are correlated with the error term is to use instrumental variables (IV). In order to employ IV one needs to augment one's model by adding variables which extend the model backwards from the independent variables of interest. Specifically, one needs to find variables, based on plausible theoretical assumptions, which affect the one or more independent variables of interest but do not directly affect the outcome. Whether a model is identified or, preferably, over-identified is then a function of the success the researcher has had in extending the theoretical specification of their model. No simple purely statistical solution is possible.

We propose an analogous method for dealing with the problem of identification of APC models. Whereas IV involves adding variables that extend one's theoretical 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 the case of IV, whether a particular model is identified depends on the theoretical richness of the specification. Thus, as with IV, in some cases our approach will work and others it will not. 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 are linearly dependent. In particular, we demonstrate how his approach can be used to develop a general strategy for identifying APC models. This approach provides a number of different strategies for identifying APC models not previously recognized in the literature. Furthermore, we show that in certain circumstances, various model specification tests are available.

Most, if not all of the previous methods for the identification of APC models can be formulated within our approach. It is also the case that, as in previous work, our approach to identification involves imposing parameter restrictions, though in many cases the restrictions involved may be implicit and may be quite a bit more complicated than those previously considered. This equivalence is due to the mathematical requirement that restrictions are needed in order to identify any APC model. 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 examining parameter restrictions. This leads to a distinctly different and new way of thinking about identification.

In the next section of the paper we briefly discuss previous research. Following this, we discuss Pearl's three different criteria for identifying causal effects. We then discuss how APC models can be identified using a mechanism based approach. Following this we examine different types of APC models. We then show how Pearl's front-door criteria can be used to deal with unobserved variables. The subsequent section examines various specification tests. We then present an empirical example.

Limitations of Previous Research

Identification of an APC model is only an issue if it is believed that Age, Period, and Cohort each potentially affect an outcome. If only two of these three variables are thought to affect the outcome then identification is not a problem since linear dependence only occurs if all three variables are present. A large number of papers in fact achieve identification by simply assuming that only two of the three variables in an APC model affect the outcome. This is a very strong theoretical assumption which may or may not be justified in particular circumstances.

As noted above, the discussion of identification within the technical APC literature has focused on placing restrictions on parameters in order to identify a model. This is typically done in two ways. First, as suggested in Mason et al. (1973), some set of parameters may be constrained to be equal. For example, it may be assumed based on some theoretical argument that the parameters associated with two periods should be constrained to be equal. This strategy has been used by Mason et al. (1973), Knoke and Hout (1974), Harding and Jencks (2003) and others. More generally identification might be achieved by assuming that two age parameters, two period, or two cohort parameters are equal. 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 specifying restrictions. (See Sasaki and Suzuki 1987 for an application).

A second 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 (Fienberg and Mason 1985b, Kahn and Mason 1987), or a period effect might be restricted to be proportional to the unemployment rate. Heckman and Robb (1985a) 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. The proxy variable approach is closest to that developed in this paper. Typically, the proportionality constraint is justified by asserting that the mechanism by which the variable of interest (Age, Period, or Cohort) affects the outcome is captured by the variable used to constrain that variables effect. O'Brien (2000) provides the most advanced discussion of this strategy.

As discussed in the introduction, though imposing restrictions certainly provides a solution to identifying the APC model, existing literature expresses dissatisfaction. First, it is often difficult to find restrictions that can be theoretically justified. Second, if the restrictions are even mildly wrong, this can have a major effect on parameter estimates (Glenn 1976). Third, apart from identification issues, much research has been criticized for being atheoretical and solely interested in decomposing the relative effects of Age, Period, and Cohort.

A fourth issue is whether it makes sense to talk about Age, Period, and Cohort as having a causal effect on some outcome. From a counterfactual perspective (for reviews see, Winship and Morgan 1999, Winship and Sobel 2003), the assertion that age, period, or cohort have causal effects is highly problematic. Holland (1986) as well as others have argued that only manipulable variables can have a causal effect (for further discussion see Winship and Sobel 2003). In other words, only variables for which it is possible to potentially change an individual's value on the causal (treatment) variable should be the subject of counterfactual causal analysis. Age, Period, and Cohort are not manipulable in this sense. There is no way that an individual's age, or birth cohort can exogenously be changed. Similarly, there is no way to exogenously change the relevant period.

Often, though, the mechanisms through which they are assumed to work are manipulable. For example, we might be interested in the importance of aging because of its association various life cycle statuses such as marriage. Although we cannot change an individual's age, the counterfactual involving having a different marital status is certainly conceivable. Similarly, Period might be of interest because of its association with unemployment rates. Here it is also conceivable to think of a particular time period as having an unemployment rate than that actually observed. Finally, Cohort might be of interest because of the potential effects of cohort size on the outcome of interest. Although we cannot change a person's cohort, the counterfactual condition in which their cohort was of a different size is certainly imaginable.

How then are we to understand causal effects within APC models? What the above discussion suggests is that makes perfect sense to think of the causal effects of particular mechanisms associated with Age, Period, and Cohort when these are manipulable. Can, however, we posit that Age, Period, and Cohort variables, themselves, causally affect these mechanisms? We would argue, that at least in a counterfactual sense, it does not. Rather, Age, Period, and Cohort represent different "clocks" or "times" that underlay different social processes. As such, their relationship with specific mechanisms might better be thought of as associational as opposed to causal. Thus, a child's height increases with age, but we would not generally think it is correct to say that aging *causes* a child to grow. Given this understanding of causality, one of the goals of an APC analysis should be understood as partitioning the changes in the outcome of interest into those components of change associated with the three dimensions of time in an APC model, not that it estimates the causal effects of Age, Period, and Cohort.

The argument for the importance 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 the processes work (e.g. Reskin 2003, Hedstrom and Swedberg 1998, Sorensen 1998). This work has argued that much sociological theory is too abstract, and in order to generate testable hypotheses about particular processes, it is necessary to specify the mechanisms involved. For example, Reskin (2003) argues that in order to test for and understand discrimination, one needs to identify the mechanism by which it occurs. One cannot simply refer to gender or race differences. This paper makes a parallel argument - that in order to achieve identification of APC models, it is necessary to specify the mechanisms through which the processes of interest work. The essential point in both the theoretical literature and here are the same: in order to know what is happening in a particular context, one needs to be able to identify the mechanisms involved.

Pearl and the Identification of Causal Effects

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

In many aspects, Pearl's theory is similar to the standard theory of linear paths models of Wright (1921) and developed within sociology by Duncan (1975). It, however, 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). It is the first two of these differences that allows us to use his theory to develop a formal theory of identification of APC models. In order to maximize the accessibility of the presentation, we minimize the formality of the math and stress intuition. Doing so makes Pearl's theory and our application of it look extremely similar to the standard path analysis model. The differences, however, are considerably greater than they appear. The reader is referred to Pearl's 2000 book for a more formal presentation (also see Pearl 1999).

The general problem that Pearl is concerned with is distinguishing true causation from simple statistical association, that is, that in many situations the association between X and Y may not provide an estimate of the causal effect of X on Y, because there are one or more variables that connect X and Y through alternative pathways and thus contribute to their association. In Pearl's theory, it is assumed that all causal variables and the associated causal relations relevant to an outcome are explicitly represented in the graph. Figure 1 shows an extremely simple example where X and Y are directly connected by a "backdoor" path through Z. Pearl describes three general strategies for identifying a causal effect from a set of observed associations.

--- Figure 1 here ---

Pearl's first principle of identification is what he calls the back-door criterion. It is a generalization of regression and involves identifying a causal effect by conditioning on some set of variables. Heckman's control function approach can be understood as a special case (Heckman and Robb 1985b).

The backdoor criterion amounts to finding variables such that if they are removed from the graph (which is statistically equivalent to conditioning on these variables) all pathways between X and Y other than the direct (causal) one are eliminated. If these variables are all observed, then the effect of X on Y is identified. The effect of X on Y is simply estimated by conditioning on the variables. This might be done through regression, matching, stratification or any other conditional method.

For example, in Figure 1 the 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 Z. Deleting Z from this graph, which in Pearl's theory is statistically equivalent to controlling for Z, eliminates this pathway. As a result, the conditional association between X and Y now estimates the causal effect of X on Y. This conditioning strategy is only possible if there are *observed* variables, the removal of which from the graph eliminates all alternative pathways between X and Y. Pearl's backdoor criterion indicates what is necessary for a conditioning strategy to identify a causal effect. Although the example here is extraordinarily simple, the backdoor criterion can be used to prove identification in much more complicated situations.

Pearl's second method of identification is the standard instrumental variable approach. As in the case 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 instrumental variables is to augment the model by adding one or more variables that (1) either directly or indirectly affect X, and (2) do not directly affect Y.

Figure 2 illustrates. The instrumental variable 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 Z.

--- Figure 2 here ---

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. The front-door criterion amounts to identifying 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 the 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.

--- Figure 3 here ---

In Figure 3, we would like to estimate the effect of S on M. The covariance/correlation between S and M does not provide a consistent estimate because of the backdoor path through U. If U is observed, then the backdoor criterion shows that we can estimate the effect of S on M by conditioning on U. If U is unobserved, this strategy is not available. However, if we can consistently estimate the effect of S on C and the effect of C on M, getting estimates of "a" and "b," then we can estimate "ab." This is the core idea behind the front door criteria.

In the present case, we can estimate both "a" and "b" by a double application of the backdoor criterion (Pearl 2000). Because there are no back-door paths between S and C, we can consistently estimate the effect of S on C. There is, however, a backdoor path between C and M through S. However, by conditioning on S, we can eliminate this backdoor path, which allowing us to consistently estimate the effect of C on M.

The example that Pearl (2000) gives is of the effects of smoking (S) on mortality (M). We are worried that there may be all sorts of reasons these two variables are associated, such that the association does not give us a consistent estimate of the causal effect. However, if we believe that smoking affects mortality primarily through lung cancer (C), then we can estimate the effect of smoking on lung cancer, "a", and the effect of lung cancer on mortality "b" to get an estimate of the effect of smoking on mortality, the product of the two effects – "ab".

Note that this methodology assumes that we know *all* causal pathways connecting smoking and mortality. If we can do this, then it is possible to identify the total effect of S on M as the sum of the effects of the causal relations connecting them. This is the key assumption of this approach. If there are other pathways, and we have not specified or cannot consistently estimate them, then we will have failed to account for those pathways through which smoking (S) effects mortality (M).

Identification of 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 would have little application to sociology since it would be too hard to find the intermediary variables, the C's. The front door criterion, however, provides a powerful 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.[1]

The basic idea behind the front door criteria is to achieve identification by adding variables to one's model that are intermediate between the independent variable or variables of interest and the outcome variable. In most cases these variables would represent the mechanisms

[1] From a counterfactual perspective (Winship and Morgan 1999, Winship and Sobel 2003), the assertion that age, period, or cohort have causal effects is highly problematic. Holland (1986) as well as others have argued that only manipulable variables can have a causal effect. In other words, only variables for which it is possible to potentially change an individual's value on the causal variable should be the subject of counterfactual causal analysis. Age, Period, and Cohort are not manipulable in this sense. Often, though, the mechanisms through which they are assumed to work are manipulable. Thus an APC model does not estimate causal effects in the counterfactual sense, but rather partitions the changes in a variable to the processes that work through Age, Period, and Cohort.

through which the original independent variables affect the outcome. The hope is that although the original model is not identified, that the subcomponents of the new model will be identified, leading to the identification of the entire new augmented model. Because the augmented model contains intermediate variables, there are now additional endogenous variables besides the final outcome of interest. Associated with each endogenous variable is there an equation. In order for the overall model to be identified each equation must be separately identified. Standard identification conditions apply to each equation of the model.

The above discussion implies two conditions that are necessary, though they may not be sufficient, for identification of a model. First, any mechanism that is directly affected by Age, Period, and Cohort must either be affected by only two of these three variables or there must be some parameter restriction that will allow the equation where this mechanism variable is the dependent variable to be identified. Second, the variables measuring the mechanisms that directly affect the final outcome must also be linearly independent or there must be some parameter restriction that allows the equation predicting the final outcome to be identified. In the examples below we focus on models where identification is achieved through linear independence as opposed to parameter restrictions since we believe that models of this type are most likely to be theoretically plausible.

Consider a simple example. Assume that we have three <u>sets</u> of dummy variables with a variable for each level except for some arbitrary base group for (1) each age, DA, (2) each period, DP, and (3) each cohort, DC. These variables could be measured at whatever level of refinement desirable – years, months, days, hours. As is well known, we cannot separately estimate the effects of all three sets of variables since they are functionally interdependent as defined by the following equality: Age = Period – Cohort.

Although the relationship between Age, Period, and Cohort can be specified in terms of an exact, deterministic mathematical relationship, social scientists in general, and sociologists in particular, often think of them as representing 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. being employed, married, having children, retired, widowed, or empty nesters). Changes with respect to period would represent the effects of current condition of society – for example, if we were referring to the U.S., 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 a recession. Finally, a cohort effect could represent the effect of being born during a specific period, (Elder's [1974] work *The Children of the Great Depression* being the most famous example) or specific properties of a cohort, such as its size. Thus, the problem is that although it is easy to specify distinct social processes associated with the general processes associated with Age, Period, and Cohort, it is not possible to straight-forwardly estimate the effects of Age, Period, and Cohort because of their linear dependence.

As discussed above, one way of placing restrictions in order to identify an APC model is to assume that one or more of the respective effects of Age, Period, or Cohort are proportional to different substantive variables (Fienberg and Mason 1985b). This is equivalent to representing each of the variables in terms of some substantive variable. Heckman and Robb (1985a) term this the "proxy" variable approach. O'Brien (2000) terms this an APC- category model. As an example, the effect of period might be specified in terms of the unemployment rate or the effect of cohort might be specified in terms of cohort size. As discussed in more detail below, as long as the parameterization is not linear, or does not otherwise create strict linear dependence between the three new representations of Age, Period, and Cohort, it is possible to estimate separate age, period, and cohort effects. In other words, identification is achieved through one or more nonlinear transformations. Multiple examples of this exist in the literature (e.g. see Fienberg and Mason 1985, O'Brien and Gwartney-Gibbs 1989, O'Brien, Stockard, and Isaacson 1999, O'Brien 1989, 2000).

Now consider how this relates to Pearl's front door criterion. Let FA be a parameterization of the Age effect, and FP and FC be analogously defined. Then we could represent these relationships in terms of the diagram in Figure 4.

--- Figure 4 here ---

One way to think about the advice that one should parameterize age, period, and cohort effects in terms of different substantive variables is that this is an application of Pearl's front door criterion. In the model in Figure 4, we can "estimate" the "a" coefficients since it simply consists of the deterministic relationships between our Age, Period, and Cohort dummies. This part of the model is deterministic, and it is obviously totally parametric. A necessary condition for estimating the "b" coefficients, that is, the effects of FA, FP, and FC, is that they be linearly independent. Once we have an estimate of the "a" and "b" coefficients, we can then calculate the total age, period, and cohort effects in terms of the appropriate products. We need, however, to consider the necessary conditions for the F variables to be linearly independent.

In the most general form, Pearl's theory is nonparametric. In our context, the easiest way to carry this out would be to create a separate dummy variable (except for the base categories) for every separate value of DA, DP, DC and FA, FP, FC. If the separate values of the D variables map onto unique values of the F variables (that is, there is a one-to-one mapping), then the model will not be nonparametrically identified since we have simply reproduced the same linear dependence among the F variables that exists among the D variables. Identification, however, can be achieved in three ways.

First, if we are willing to restrict the functional form of the relationship between the F variables and the outcome Y, for example as linear, then identification can be achieved as long as the F's (in non-dummy variable form) are not linear dependent. Since they are nonlinear transformations of the D variables, this is unlikely, though not impossible (see Heckman and Robb 1985a). A second possibility is that identification is achieved if some of the separate values of each D variable map onto the same value of an F variable. In this case, the specification of a relationship between a D variable and an F variable is equivalent to restricting some set of the D dummy variables to have equivalent effects. For example, if we parameterized the period effect by the unemployment rate, and two periods had the same unemployment rate, this would be equivalent in a nonparametric model to equating the effects for the two dummy variables for the two periods. As shown by Mason et al. (1973), this is all that is needed in order to achieve identification. A third condition that allows identification, but falls outside the re-parameterization case, is when there are some individuals who have equivalent values on their D variables, but different values on their F variables. Here, nonparametric identification is also possible. We discuss this case in detail below.

Alternative Types of APC Models

The front door approach suggests that we can identity the effects of variables by introducing intermediary variables that specify the mechanism(s) by which our variables of interest affect the outcome. As in the case of instrumental variables, augmenting our initial model allows us to identify causal effects of interest. In the case of instrumental variables, the model is augmented by introducing one or more variables that affect the independent variable of concern, but that do not directly affect the outcome. In the case of the front-door criterion, identification is achieved by introducing intermediary variables that are not correlated with the unobserved variables of concern. In this paper, however, intermediary variables are introduced, not in order to deal with unobserved variables per se, but to deal with the identification problem introduced by the deterministic functional dependence between Age, Period, and Cohort. Thus, in any particular analysis, the introduction of intermediary variables can be done for two reasons. First, as discussed by Pearl, it may be done to identify causal effects where those causal effects are not identified due to unobserved variables. Second, intermediary variables may be introduced in order to deal with the problem dependence between the variables of interest associated with one or more unobserved variables.

There is nothing about the front door criterion that assumes that the intermediary variables must simply be nonlinear transformations of the independent variables of interest. In particular, the intermediary variables could be different substantive variables that specify particular causal mechanisms that relate each independent variable to the outcome of interest. If we recognize this, it is clear that there is a much richer set of models that are identified than those that have been typically been considered. To see why this is the case, consider Figure 4 again. This model contains multiple restrictions. First, it assumes that none of the D variables directly affect the outcome. This amounts to three restrictions. We will use this fact later as the basis for developing a general misspecification test. Second, each D variable is assumed to affect only one F variable. This amounts to six additional restrictions. Thus the model in Figure 4 has a total of nine restrictions. As discussed above, only one restriction is needed to identify an APC model. Thus, more general models that do not contain these restrictions can be considered and, because the model is over-identified, its fit can be tested.

--- Figure 5 here ---

Consider Figure 5. This model is fully estimable as long as one of the three conditions discussed above hold. The effects of DA and DP on T and similarly the effects of DP and DC 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.

Note that this model involves fewer restrictions than the model in Figure 4. 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.

Unobserved Variables

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, perhaps the biggest 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 have failed to estimate the total effect of one or more of these variables on the outcome. Our approach makes the very strong assumption that we have identified all the mechanisms through which Age, Period, and Cohort work. As we discuss in more detail below, however, this condition can be relaxed. All that is necessary to identify the model is that it be the case for one of the APC variables that all the mechanisms by which it affects the outcome be identified. When this is the case, the effects of the other two APC variables can be controlled for by simply including them in the equation predicting the outcome.

We now consider the problem of unspecified mechanisms more explicitly. Doing so demonstrates the power and limitations of Pearl's identification theory, particularly the front door criterion. Building on this analysis, in the following section we examine two general specification tests that assess the overall fit of the model. Since a model may be misspecified in a variety of ways, these tests are particularly helpful in determining whether a specification is appropriate.

Consider Figure 6, which is identical to Figure 4 except that there is an additional path connecting DA and Y through an unobserved variable UA. UA should be thought of as an unspecified or unobserved mechanism. 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 criteria states that if we can consistently estimate the "a" and "b" coefficients, then we can consistently estimate the total effects of Age, Period, and Cohort as their appropriate products. For the moment assume that the estimation of the "a" coefficients is unproblematic. Also assume that the F functions are <u>not</u> deterministic functions of each other.

--- Figure 6 here ---

In order for the "b" 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 UA) connecting each F variable and Y.

Consider the problem of estimating the effect of FC on Y, "b_3." There are backdoor paths between FC and Y through FP, FA, and UA via the D variables. If there were no unobserved UA variable, as in Figure 4, then "b_3" could be consistently estimated by simply conditioning on FA and FP, say, by using a regression model (as long as they are not deterministic functions of each other). Above, we discussed the necessary conditions for this to be true.

In contrast, in Figure 6, conditioning on FA and FP still leaves the backdoor path FA-DA-UA-Y. This path, however, could be identified by conditioning on DA. Since by assumption DA, FP, FC are not exact functions of each other (which would be the case in most empirical applications), this effect is identified.

Now consider the problem of estimating "b_1." If FA is a deterministic function of DA, then it will not be possible to estimate "b_1" conditioning on DA. There will be an interdependence problem. Let's say, however, that there is variation in FA independent of DA. This would generally be true for a variable like marital status. Because there is independent variation in FA, it will be possible to estimate "b_1" by conditioning on DA. Note that there is no need to condition on either FP or FC. Conditioning on DA breaks down all backdoor paths between FA and Y.

Assume, however, that the model is a bit more complicated and that FA is also affected by DC. In this case there would now be the backdoor path FA-DC-FC-Y between FA and Y. Here, we would need to condition on FC or DC as well as DA in order to break all backdoor paths between FA and Y.

The example of "marital status" shows that there is an additional identification strategy in APC models. Above, we noted that the variable paramaterization method, Heckman and Robb's proxy variable approach, and O'Brien's APC-category 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, Y. The marital status example demonstrates that in the case in which an intermediary variable contains variation independent of the D variables, then its affect can also be identified by conditioning on the D variables since the variable is not an exact deterministic function of the D variables.

Finally, consider the problem of estimating the total effect of DA on Y. This is equal to $(a_0\ b_0) + (a_1\ b_1)$. Logically, there is no reason that we cannot simply drop UA and FA from the graph in Figure 6 and draw a single line between DA 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 DA and Y through both DP and DC. Conditioning on both DP and DC is not possible because of the dependence problem between these three variables. We could, however, break these backdoor paths by conditioning on any of the following pairs of variables: DP and FC; FP and DC; or FP and FC. Note that the first of two of these pairs demonstrates how in order to identify the total causal effects of Age, Period, or Cohort we only need to have specified the complete set of mechanisms associated with one of these three variables. Thus, in this example it is possible to identify all three effects if either FP or FC are unobserved, though identification is not possible if this is true for both. O'Brien's (2000) APC-characteristics models embody this insight.

Obviously, numerous other models could be examined. The point here is two-fold. The first is to show how Pearl's backdoor and front door criteria work and to show their power. The second is to show how conditioning on one or more of the D variables allows one to deal with unobserved variables. This is one of oldest insights in the APC literature. Here, we provide a formal theory for its application.

Specification Tests

The fact that it is possible to specify more complicated APC models means that it is possible to carry out various specification tests. We discuss three here. First, when one model is nested within another, we can use standard F tests or Chi-square likelihood ratio tests to compare the models. For example, in Figure 5 above, we can test whether in fact Cohort affects T. What this example shows is that by using more elaborate models, the researcher can test whether specific mechanisms mediate the effect of Age, Period, or Cohort.

Second, under certain conditions, we can perform more general specification tests to determine whether we have fully captured the effects of age, period, and cohort in our model. As Feinberg and Mason (1979, 1985) point out, it is possible to test the fit of a model even though all the parameters of the model may not be identified. Consider a test that assesses whether there are other separate functions of Age, Period, and Cohort that affect the outcome that have been omitted from the model in Figure 4. This test consists of the following steps:

Step 1: Regress Y on FA, FP, and FC and calculate R-square.

Step 2: Regress Y on FA, FP, FC, DA, and DP and calculate R-square.

Step 3: Perform an F-test.

If we have correctly specified the separate effects of Age, Period, and Cohort, then to within sampling error, the two R-squares should be the same. This can be formally tested using an F-test since the two models are nested within each other. If they are not equal, then this is evidence that there are other causal pathways between Age, Period, Cohort and the outcome variable that have not been accounted for.

A more general form of this test is to include not just separate dummied Age and Period effects in step two but, in addition, to include all possible interactions between the two variables. In this case we are fitting the saturated model. That is, implicitly we are testing whether there are any variables correlated with any function of Age, Period, or Cohort that have been omitted from the model.

Both of these last tests should be understood as general specification tests. Specifically, they can be thought of as Hausman type tests (1978). In both cases, we are testing whether arbitrary functions of Age, Period, and Cohort have predictive power with respect to the outcome over and above the mechanisms specified in the model. As such, we may fail to accept a specific model for a variety of reasons. As already noted, one possibility is that there are intermediary variables that have been omitted from the model. A second possibility is that the intermediary variables have measurement error. A third possibility is that there are other variables correlated with or that affect Age, Period, and Cohort, that also affect the outcome.

Empirical Example

To illustrate these ideas, we conduct a basic analysis of the effects of Age, Period, and Cohort on political alienation. Following Kahn and Mason (1987), we use data from white males surveyed by the National Election Surveys for presidential election years.[2] Here, political alienation is measured by whether the respondent agrees or disagrees with the statement: "I don't think public officials care much what people like me think." Those who agree are coded one, and those who disagree are coded zero. Other variables are described in Table 1. We restrict our analyses to those aged 30 to 60 in 1956, 1960, 1964, 1968, 1976 and 1980 who have no missing data on any of our variables, leaving 2041 cases.[3]

--- Table 1 here ---

For expository clarity, throughout this example we use OLS for all models even though

[2] Data are from the National Election Studies Cumulative Data File, 1948-2000, ICPSR No. 8475 (Sapiro, Rosenstone, and the National Election Studies 2002).
[3] Kahn and Mason (1987) also include those surveyed in 1952 and 1972, but an important variable for our example, Number of Kids, was not available in those years.

OLS is not appropriate for many of the models because the dependent variables are binary, ordinal categories, or nominal categories. Probit, Ordered Probit, Logit, or Multinomial Logit models would be more appropriate for many of the models. However, since we provide this example for pedagogical rather than substantive purposes and want it to be accessible to the widest possible audience, we use OLS regression. Results are similar when more appropriate models are used in place of OLS.

--- Figure 7 here ---

Figure 7 shows a simple model of the relationships between Age, Period, Cohort, and Political Alienation (hereafter, PA). Analogous to Figure 4, it specifies a single intervening mechanism for each variable of interest: Age, Period, and Cohort. (For now, ignore the dashed arrow labeled a_3). This model makes a number of assumptions. First, it assumes that the effect of Period on PA operates entirely through the unemployment rate, the effect of Age on PA operates entirely though the number of kids one has, and the effect of cohort on PA operates entirely through cohort size. Second, it assumes that all causal relationships between variables in the diagram are indicated by arrows, i.e. that Cohort does not affect the number of kids one has, and so on. If these assumptions are correct (and we have avoided other common problems like misspecification of functional form, measurement error, etc.), we can easily estimate the effects of Age, Period, and Cohort on PA. We simply estimate a model predicting PA with the Unemployment Rate, Number of Kids, and Cohort Size to estimate b_1, b_2, and b_3, estimate a model predicting number of kids with Age to estimate a_2, and estimate a model predicting cohort size with cohort to estimate a_4.[4] Because the relationship between Period and Unemployment Rate is deterministic, we already know a_1 (see Table 2). The effects of Period, Age, and Cohort on PA are then (a_1b_1), (a_2b_2), and (a_4b_3), respectively.

[4] This example ignores the known relationship between period and cohort size for purposes of exposition. In our data, cohort size is actually completely determined by period, cohort, and their interaction.

--- Table 2 here ---

Table 2 indicates the relationship between each level of Period and Cohort and their proxy variables. Since the relationships between Period and Unemployment Rate, between Period and Watergate, and between Period and Republican President are deterministic, we do not discuss them further. Note that Unemployment Rate is highly nonlinearly related to Period. Table 3 reports parameter estimates.

--- Table 3 here ---

In column 1 of Table 3 we see that Age dummies have large and statistically significant effects on Number of Kids. R^2 for this model is 0.122. Column 4 shows that the Unemployment Rate (b = 0.027, se = 0.007) and Cohort Size (b = -0.033, se = 0.009) but not Number of Kids (b = 0.004, se = 0.007) have modest but significant effects on Political Alienation. The R^2 for this model is 0.149.

We can test whether the model is more generally misspecified using the two general tests described in the previous section. Table 4 reports on the model specification tests used with different variants of our empirical example.

--- Table 4 here ---

In the present case, we first estimate a model predicting PA with the Unemployment Rate, Number of Kids, and Cohort Size. Then we estimate a second model which also includes sets of dummy variables for Period and Age. The F test which compares these two models tests whether Period and Age (and also Cohort, since it is a deterministic function of them) add any explanatory power. If they do, then we have not accounted for all the intervening variables between Age, Period, and Cohort and our outcome, PA. In this case, the test rejects the null hypothesis that the R^2s of the two models are the same. As reported in row 1 of Table 4, the F-test statistic is 7.55 with 11 and 2026 degrees of freedom (one of the period dummies was dropped due to collinearity), which has $p < .001$. Adding in the full set of interactions into the model gives, as shown in row 2 of Table 4, an F-statistic of 2.69 with 45 and 1992 degrees of freedom, which is also significant at the $p < .001$ level. Interestingly, the difference between the two models, which is a test for whether the interaction terms are significant, gives an F statistic of 1.11 with 34 and 1992 degrees of freedom, which is not statistically significant ($p = 0.304$; see row 3 of Table 4).[5] This suggests that there are missing mechanisms in our model, as certainly there are, but that the linear specification is appropriate.

As noted previously, a general specification test can fail for a number of reasons, all of which indicate we have not correctly specified the model. In addition to omitted intervening variables, it might also be that we have measurement error in our intervening variables, that we do not have the correct functional form or that there are uncontrolled for/unmeasured variables that are correlated with our independent variables that affect our outcome.

It is also possible to test, in a more limited way, whether additional arrows between variables are necessary in the model. For example, we might wonder whether Cohort affects Number of Kids, which can be represented in Figure 7 with the additional arrow labeled a_3. Column 2 in Table 3 shows a new model for Number of Kids, which is now predicted by both Age and Cohort. These estimates suggest that Cohort does affect Number of Kids, as a number of the Cohort dummies are large and statistically significant. An F test comparing models 1 and 2 shows that the two models are significantly different. The test statistic is 5.00 with 13 and 2020 degrees of freedom, which has $p < 0.001$ (see row 4 of Table 4).

--- Figure 8 here ---

[5] Where we using a more appropriate model for PA, such as a maximum likelihood Probit model, then we could simply substitute likelihood ratio tests for the F tests.

As discussed above, we need not specify a single intervening variable for each variable A, P, and C. As long as each intervening variable is affected by only two of the three, we should be able to estimate the model. Figure 8 is analogous to Figure 5, and specifies two intervening variables, each related to two of the three exogenous variables. The effect of A on PA operates through both employed and church attendance and is estimated as $(a_2b_1)+(a_3b_2)$. The effect of P on PA operates only through employed and is estimated as a_1b_1. And the effect of C on PA operates only through church attendance and is estimated as a_4b_2.

--- Table 5 here ---

Table 5 reports the estimates for this model. The first column reports the effects of Period and Age on Employment. The effects are small and in some cases nonsignificant. The R^2 for this equation is 0.041. Column 2 reports the effects of Age and Cohort on Church Attendance. The effects of Age are small and but significant for most of the dummy variables. The Cohort effects are often large and nearly always significant. As reported in column 3, the effects of employment (b = -0.162, se = 0.044) and Church Attendance (b = -0.055, se = 0.009) on PA are large and significant.

Clearly, however, this is an unrealistic model in that it does not come close to representing all of the potential pathways through which A, P, and C can affect PA. This theoretical hunch is confirmed by our specification test. As reported in row 5 of Table 4, when we compare a model of PA which includes only employed and church attendance with one which also includes Period and Age dummies, the F test rejects the null hypothesis that the two models are the same. The F statistic is 8.13 with 12 and 2026 degrees of freedom, which has $p < .001$. The test with interactions, reported in row 6, yields an F statistic of 2.78 with 47 and 1991 degrees of freedom. Again the difference between these models is not significant (F = 0.94 with 35 and 1991 degrees of freedom, p = 0.569; row 7 of Table 4), indicating that interactions are not needed.

--- Figure 9 here ---

Figure 9 represents what we see as a more realistic model of the relationships between A, P, and C and the intervening variables and PA. PA is directly affected by Watergate, Republican President, Employed, Education, Number of Kids, and Church Attendance. There are two "stages" of intervening variables in this model, since some of the variables that directly affect PA are not directly affected by A, P, or C. Further intervening variables include Cohort Size, Unemployment Rate, and Marital Status. The model is complicated by the fact that some of the variables directly affecting PA are also intervening variables for other variables. For example, Number of Kids mediates the relationship between A and Church Attendance and also directly affects PA.

Table 6 provides estimates for the different equations represented by this model. Space limitations prevent us from discussing all the individual coefficients. Focusing on just the equation for Political Alienation we see that Employment, Church Attendance, Education, a Republican President, and Watergate all have substantial and statistically significant effects. Number of Kids, however, does not. The R^2 for this model is a modest but reasonable 0.114.

--- Table 6 here –-

How well does this richer model fare in representing the intervening variables? We can use our two specification tests to find out. First, we estimate a model which predicts PA with those variables directly affecting it: Watergate, Republican President, Employed, Education, Church Attendance, and Number of Kids (see Column 6 of Table 6). This model is nested within a second model which also includes Age and Period Dummy variables (model not shown). Comparing their R^2s is our first specification test. When we perform this test, we find that we cannot reject the null hypothesis that the two models have the same R^2 (From row 8 of Table 4: F statistic of 1.47 with 10 and 2024 degrees of freedom has a p-value of 0.144). This is evidence that there are no variables between A, P, or C and PA that are not represented in our model. We also performed the stronger test by including all interactions. Here we get an F statistic of 1.01 with 45 and 1989 degrees of freedom, which has a p-value of 0.463 (see row 9 of Table 4). Again, the Period by Age interactions do not seem to be necessary. Comparing a model with the interactions to one without generates an F statistic of 0.87 with 35 and 1989 degrees of freedom, which is not statistically significant ($p = 0.681$; see row 10 of Table 4).

However, since there are multiple "stages" in our model, we can and should perform the same specification test for each of the stochastic intervening variables. These tests of intervening variables are important. If we do not have all the connections between each of these variables and A, P, and C, then we will not be able to calculate the total effects of A, P, and C on PA by summing the paths. Some of the paths will be missing.

For example, the model in Figure 9 shows that Church attendance is determined by Period, Number of Kids, and Marital Status and that there are no direct connections between either C and Church Attendance or A and Church Attendance. We can test this with the specification tests. The first model predicts Church Attendance with Cohort Size, Marital Status and dummies for each Period (see Column 4 of Table 6). The second model adds our Age dummies (model not shown). According to our test, the Age variables do add to the Church Attendance model. The F statistic is 4.99 with 12 and 2024 degrees of freedom which has a p-value of less than 0.001 (see row 11 of Table 4). Row 12 of Table 4 shows the full identification test, including Age-Period interactions as well. This test too is significant ($F = 1.91$ with 47 and 1989 degrees of freedom, $p < 0.001$). These results suggest that we have not captured all the mechanisms connecting A, P, and C and Church Attendance with our model.

We can do the basic test only for variables that are affected by only one of Age, Period, or Cohort. If a variable is directly affected by any two of the three, there is no larger model in which the model is nested that can be estimated. We can, however, test for interactions. When we compare a model for Education with Age and Cohort to one that also includes (Age)x(Cohort) interactions, the models are not significantly different (F statistic of 0.68 with 27 and 1993 degrees of freedom, p = 0.794; see row 13 of Table 4).

Unfortunately, as shown in Table 4, Employed (rows 14 and 15) and Number of Kids (rows 16 and 17) both fail both their specification tests. For marital status we estimate separate models comparing each of single, divorced, and widowed to married. These models are shown in Column 6 of Table 6. As shown in Table 4, the comparison between married and divorced also fails both specification tests, though the other Marital Status comparisons pass (see rows 18 to 23 in Table 4).

These results mean we have failed to specify at least one of the paths between A, P, or C and each of the intervening variables. There are no more variables in the data that will help us to fill out these pathways. A different data set is required, but we have illustrated how to use intervening variables to achieve and test model identification in the presence of functional dependence between variables.

Note that we do not need to, and indeed cannot, perform the specification test with the non-stochastic (deterministic) intervening variables: Watergate, Republican President, and Unemployment Rate. These three variables are "mechanically" related to Period, such that once we know the period, we know the value that these variables take, by construction. There is no need to perform the test because these variables have no variance once Period is known.

Although our model has not passed all specification tests, for illustrative reasons we construct the total effect of Age, Period, and Cohort from the model in Figure 9. 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 7 shows an example calculation comparing those in the 1936-1939 cohort surveyed in 1976 with those in the 1908-1911 cohort surveyed in 1960. 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.

To calculate the total effect of either Age, Period, or Cohort on PA, we must first list all the paths through which each factor is connected to the outcome in Figure 9. The leftmost column of Table 7 lists these paths for Period, Cohort, and Age. The effect associated with each path is the product of the coefficients for each segment of the path. These can be read from the estimates in Table 6 for the stochastic variables and from Table 2 for the variables that are deterministically related to Period. The middle column of Table 7 shows the calculation for each path, and the rightmost column shows the results of the calculation. Finally, to get the total effect, we sum all the paths. The total Period effect is 0.2611, the total Cohort effect is 0.0227, and the total Age effect is -0.0252. The grand total predicted difference between the two groups based on our model is 0.2586. This means that about 26 percentage points more respondents born 1936-1939 and surveyed in 1976 were politically alienated than respondents born 1908-1911 and surveyed in 1960. Interestingly, the period effect dominates the total difference for this example, and within the period effect, most of the difference is due to Watergate and the party of the President.

For pedagogical purposes, we have used OLS models to simplify the presentation of this example and make is accessible to a wider audience. This also simplified considerably the calculations of total effects just discussed. Because the models were estimated with OLS, we could use the raw coefficients from the models to calculate the products that capture the contribution of each of the paths. When the linearity and additivity assumptions of OLS are not appropriate or one uses models for discrete data such as Logit or Probit models, raw coefficients must be first transformed to a common metric before paths are calculated. For further details, we refer the reader to previous work on path analysis for discrete data and on nonlinear and non-additive models such as Stolzenberg (1979), Fox (1980, 1985), Winship and Mare (1983, 1984), and Xie (1989).

Conclusion

Although there is a large literature on the identification of APC models, to date, it has not provided a fully satisfactory solution. Most importantly, researchers have complained about the lack of connection between theory and actual analyses. In this paper we have outlined a general approach to estimating models with over-determined outcomes based on Pearl's front-door criterion. Our approach rests on theoretically identifying the mechanisms by which Age, Period, and Cohort affect the outcome. By including theoretically specified variables through which Age, Period, and Cohort work, identification will generally be possible. In addition, the model developed will often be sufficiently rich to allow a variety of specification tests to be carried out. We have illustrated our approach by providing an extended analysis of the determinants of political alienation. In developing our approach, our hope has been that it will engender a new interest in APC models, particularly the analysis of more sophisticated models than have been considered in the past. The virtues of more sophisticated models are to be found both in providing models that are more theoretically compelling and complete and in providing the potential for a richer analytical approach.

References

Brickman, J. D. and C.H. Weiss. 2000. "Theory Based Evaluation in Practice." *Evaluation Review* 24: 407-31.

Chen, H. and Peter.H. Rossi. 1987. "The theory-driven approach to validity." *Evaluation and Program Planning* 10: 95-103.

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.

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

Fienberg, Stephen E. and William M. Mason. 1985. "Specification and Implementation of Age, Period, and Cohort Models." Pp. 45-88 in William. M. Mason and Stephen. E. Fienberg, editors. *Cohort Analysis in Social Research*. New York: Springer-Verlag.

Fox, John. 1980. "Effect Analysis in Structural Equation Models." *Sociological Methods and Research* 9:3-28.

Fox, John. 1985. "Effect Analysis in Structural Equation Models 2: Calculation of Specific Indirect Effects." *Sociological Methods and Research* 14:81-95.

Glenn, Norval D. 1981. "The Utility and Logic of Cohort Analysis." *The Journal of Applied Behavioral Science,* 2: 247-257.

Glenn, Norval D. 1976. "Cohort Analysts' Futile Quest: Statistical Attempts to Separate Age, Period and Cohort Effects." *American Sociological Review* 41: 900-904.

Harding, David J. and Christopher Jencks. 2003. "Changing Attitudes Toward Pre-Marital Sex: Cohort, Period, and Aging Effects." *Public Opinion Quarterly* 67: 211-226.

Hausman, J. A. 1978. "Specification Tests in Econometrics." *Econometrica* 46: 1251-1271.

Heckman, James, and Richard Robb. 1985a. "Using Longitudinal Data to Estimate Age, Period, and Cohort Effects in Earnings Equations." in W. M. Mason and S. E. Fienberg. Editors. *Cohort Analysis in Social Research*. New York: Springer-Verlag: 137-150.

____________________________. 1985b. Alternative methods for evaluating the impact of interventions. In *Longitudinal analysis of labor market data*, eds. J.J. Heckman, B Singer, pp. 156-245. Cambridge: Cambridge University Press.

Hedstrom, Peter and Richard Swedberg, eds. 1998. *Social Mechanisms: An Analytical Approach to Social Theory*. Cambridge: Cambridge University Press.

Holland, Paul W. 1986. "Statistics and Causal Inference (with comments)." *Journal of the American Statistical Association* 81: 945-970.

Issac, Larry W., Debra A. Street, and Stan J. Knapp. 1994. "Analyzing Historical Contingency with Formal Methods." *Sociological Research and Methodology*, 23: 114-41.

Kahn, Joan R. and William M. Mason. 1987. "Political Alienation, Cohort Size, and the Easterlin Hypothesis." *American Sociological Review* 52:155-169.

Knoke, David and Michael Hout. 1974. "Social and Demographic Factors in American Political Party Affiliations, 1952-1972." *American Sociological Review* 39: 700-713.

Mahoney, J. 1999. "Nominal, Ordinal, and Narrative Appraisal in Macrocausal Analysis." *American Journal of Sociology*, 104: 1154-96.

Mahoney, J. 2000. "Strategies of Causal Inference in Small N Analysis." *Sociological Methods and Research,* 28: 387-424.

Mason, Karen O., H.H. Winsborough, William M. Mason, 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, editors. 1985a. *Cohort Analysis in Social Research*. New York: Springer-Verlag.

Mason, William M. and Stephen E. Fienberg. 1985b. "Introduction: Beyond the Identification Problem." Pp. 1-8 in William. M. Mason and Stephen. E. Fienberg, editors. *Cohort Analysis in Social Research*. New York: Springer-Verlag.

Mills, John Stuart. [1843] 1974. *A System of Logic.* Reprint, Toronto, Canada: University of Toronto Press.

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. 1989. "Relative Cohort Size and Age-Specific Crime Rates: An Age-Period-Relative Cohort Size Model." *Criminology* 27: 57-78.

O'Brien, Robert M. 2000. "Age Period Cohort Characteristic Models." *Social Science Research* 29: 123-139.

O'Brien, Robert M. and Patricia A. Gwartney-Gibbs 1989. "Relative Cohort Size and Political Alienation: Three Methodological Issues and a Replication Supporting the Easterlin Hypothesis." *American Sociological Review* 54: 476-480.

O'Brien, Robert M, Jean Stockard, and Lynne Isaacson 1999. "The Enduring Effects of Cohort Characteristics on Age-Specific Homicide Rates, 1960-1995." *American Journal of Sociology* 104: 1061-1095.

Pearl, Judea. 1999. "Graphs, Structural Models, and Causality," in *Computation, Causation, & Discovery*, edited by Clark Glymour and Gregory F. Cooper. Cambridge: MIT Press.

Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference.* Cambridge: Cambridge University Press.

Ragin, Charles. 1987. *The Comparative Method: Moving Beyond Qualitative and Quantitative Strategies.* Berkeley: University of California Press.

Reskin, Barbara. 2003. "Including Mechanisms in Our Models of Ascriptive Inequality." *American Sociological Review* 68: 1-21.

Sapiro, Virginia, Steven J. Rosenstone, and the National Election Studies. 2002. *American National Election Studies Cumulative Data File, 1948-2000.* [computer file]. 11th ICPSR Version. 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.

Sorenson, Aage. 1998. "Theoretical Mechanisms and the Empirical Study of Social Processes." Pp. 238-266 in Peter Hedstrom and Richard Swedberg, eds. *Social Mechanisms: An Analytical Approach to Social Theory.* Cambridge: Cambridge University Press.

Stolzenberg, Ross M. 1979. "The Measurement and Decomposition of Causal Effects for Nonadditive Models." Pp. 459-99 in *Sociological Methodology 1980*, edited by Karl F. Schuessler. San Francisco: Jossey-Bass.

Weiss, C. H. 1998. *Evaluation Methods for Studying Programs and Policies.* Upper Saddle River, NJ: Prentice Hall.

Winship, Christopher and Robert D. Mare. 1983. "Structural Equations and Path Analysis for Discrete Data." *American Journal of Sociology* 89: 54-110.

Winship, Christopher and Robert D. Mare. 1984. "Regression Models with Ordinal Variables." *American Sociological Review* 49: 512-525.

Winship, Christopher and Stephen L. Morgan. 1999. "The Estimation of Causal Effects from Obervational Data." *Annual Review of Sociology* 25: 659-707.

Winship, Christopher and Michael Sobel. 2003. "Causal Inference in Sociological Studies," in

Melissa Hardy and Alan Bryman, eds. *Handbook of Data Analysis*. Sage Publications.

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 and Research* 17: 325-352.

Figure 1

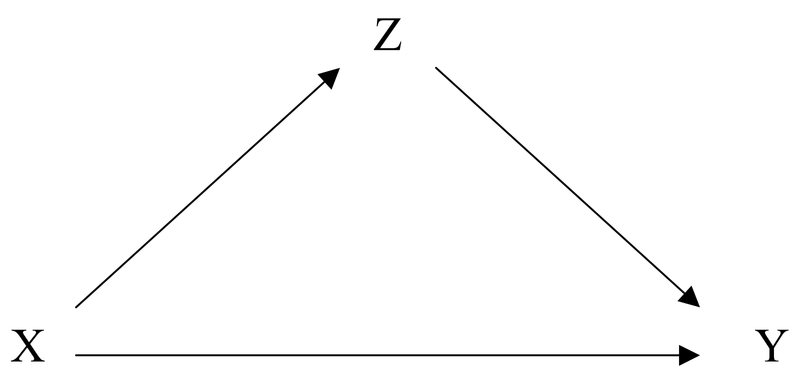

Figure 2

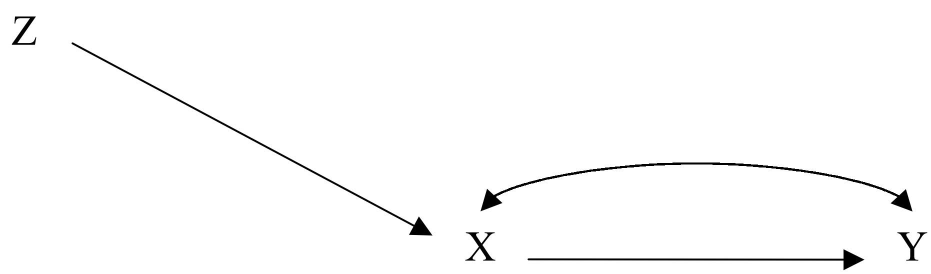

Figure 3

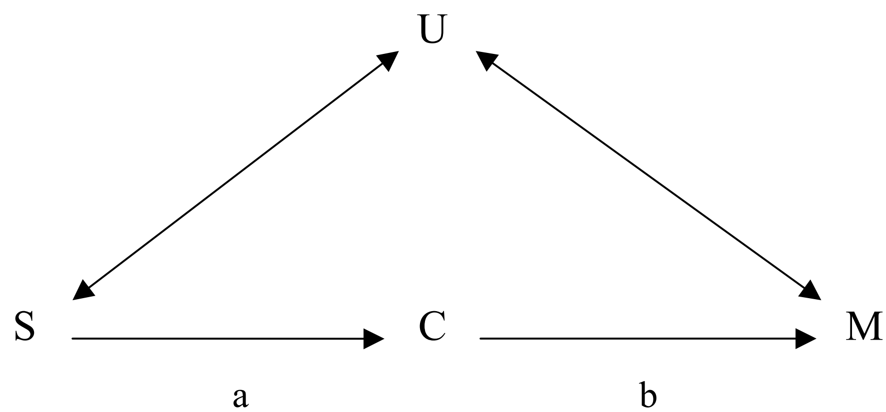

Figure 4

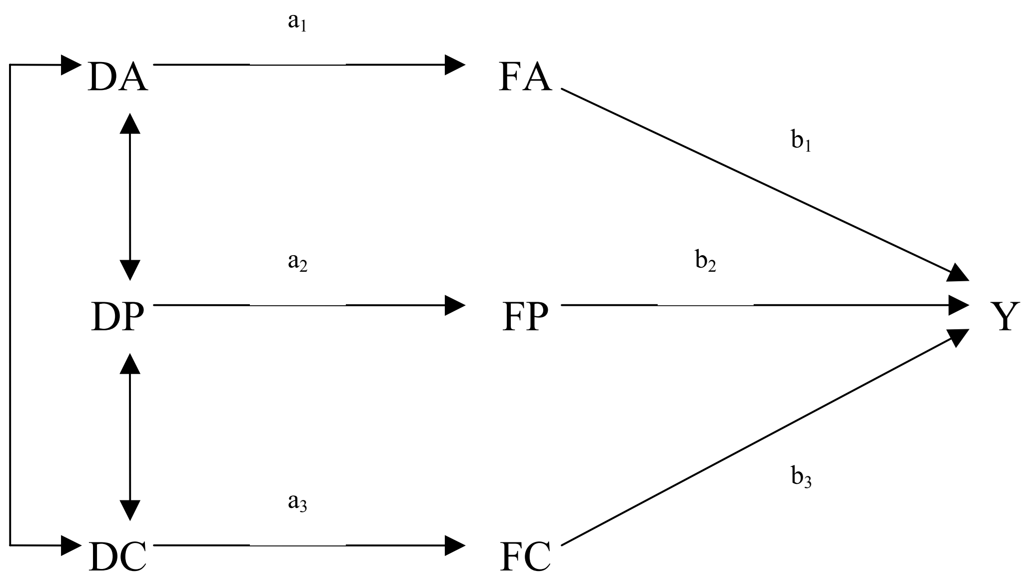

Figure 5

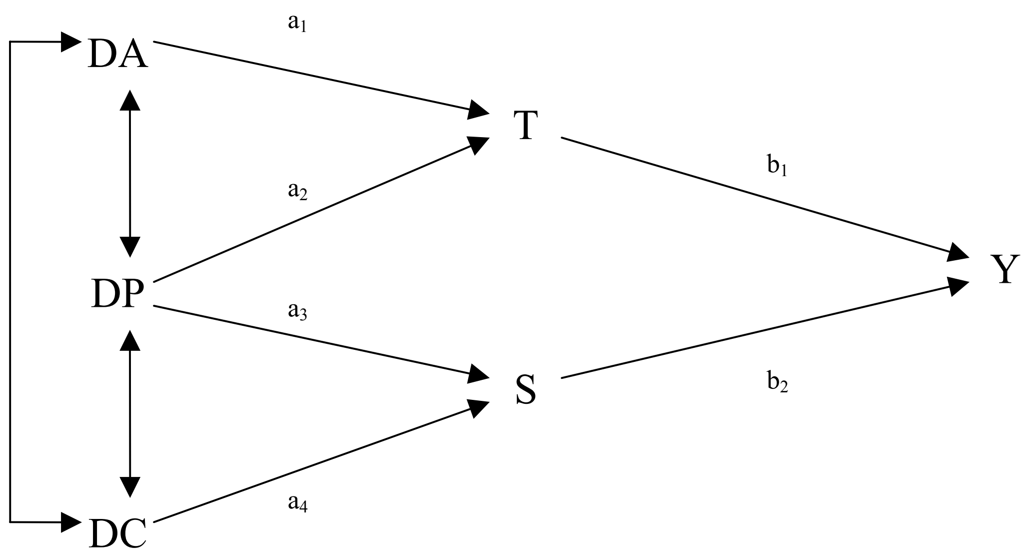

Figure 6

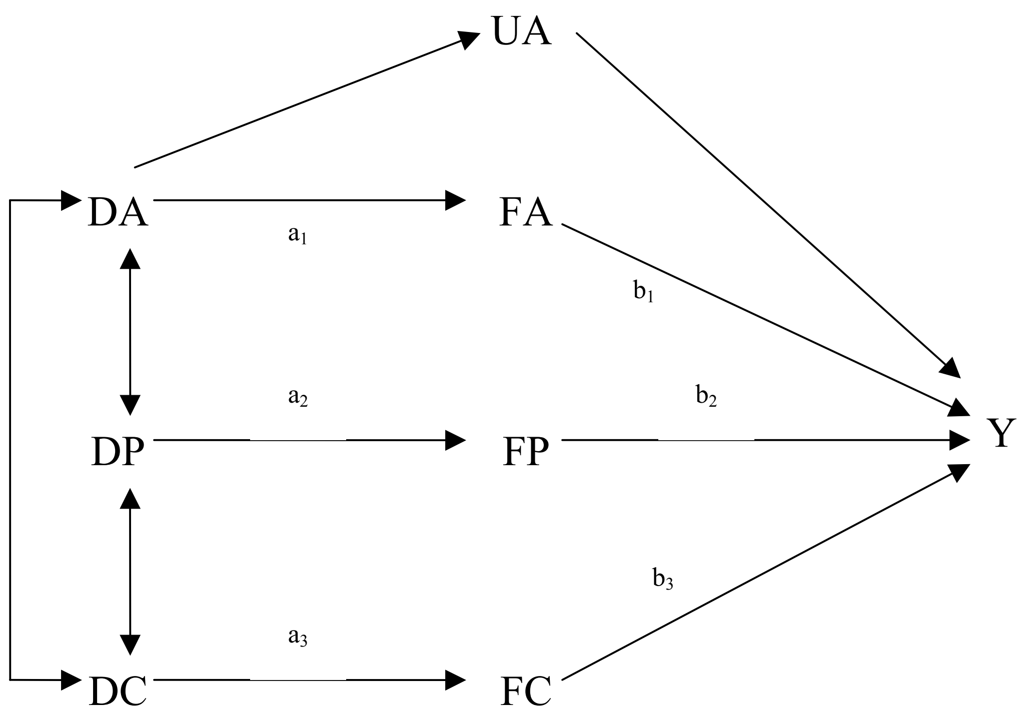

Figure 7

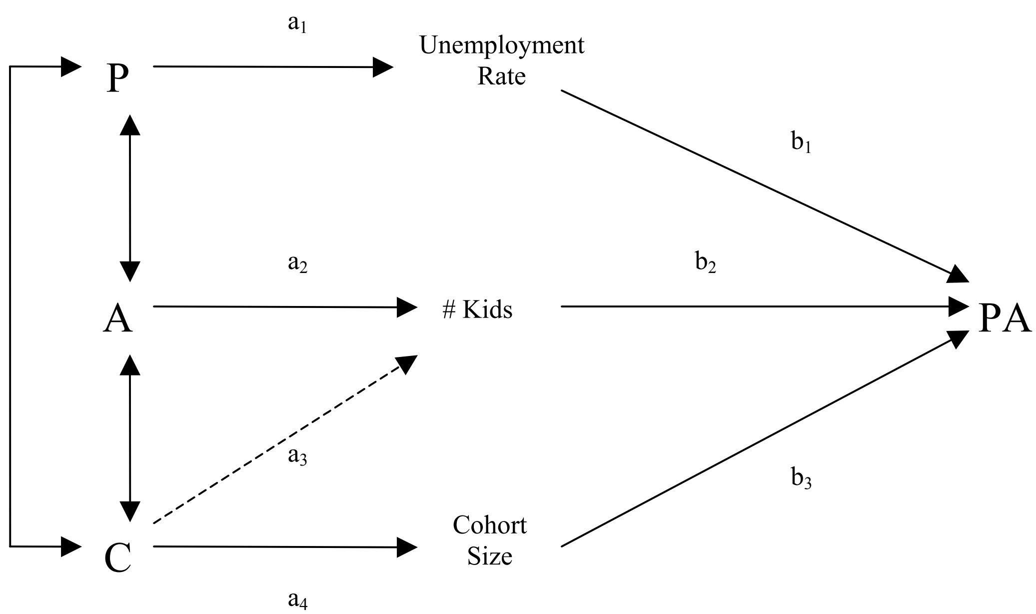

Figure 8

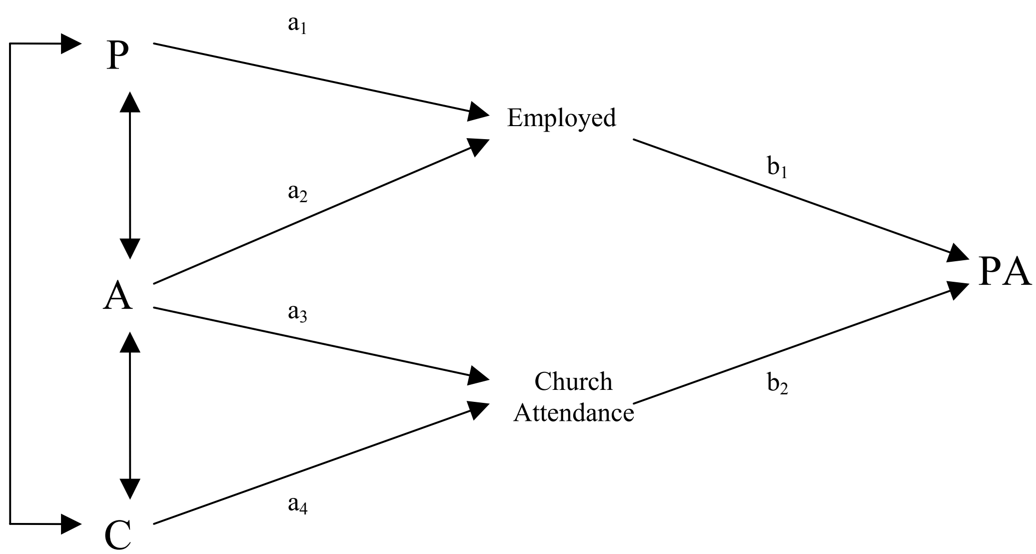

Figure 9

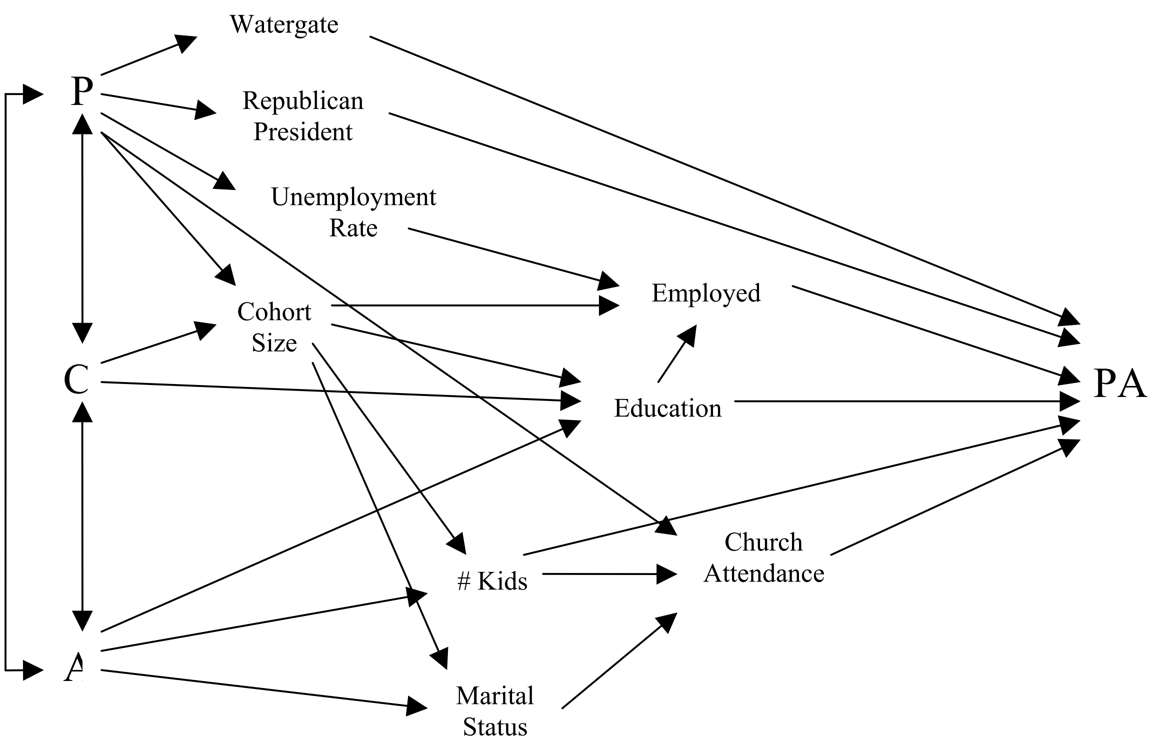
Watergate
P
Republican President
Unemployment Rate
Employed
Cohort Size
C
Education
PA
Church Attendance
Kids
A
Marital Status

Table 1: Variable Descriptions for Political Alienation Example

Name	Description	Source
Political Alienation (PA)	Agree with statement in text.	NES (VCF0609)
Age (A)	30 to 60	NES (VCF0101)
Period (P)	Years: 1956, 1960, 1964, 1968, 1976, 1980	NES (VCF0004)
Cohort (C)	Categories based on 4-year birth year intervals	Kahn and Mason (1987)
Cohort Size	Percent of US White Males in Cohort in Year	Table 1, Kahn and Mason (1987)
Unemployment Rate	Unemployment rate for US Males age 20+ in November	CPS
Republican President	1 in years in which sitting President is Republican (1956, 1960)	
Watergate	1 in years after Watergate Scandal broke (1976, 1980)	
Employed	1 if currently employed	NES (VCF0118)
Education	Years recoded from categorical variable	NES (VCF0110)
Number of Kids	Number of people under 18 in household (topcoded at 8)	NES (VCF0138)
Marital Status	1 = Married 2 = Single 3 = Divorced/Separated 4 = Widowed	NES (VCF0147)
Church Attendance	1 = Never 2 = Seldom 3 = Often 4 = Regularly	NES (VCF0130, VCF0131)

Table 2: Deterministic Relationships

Period	Unemployment Rate	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

Table 3: OLS Estimates for Models in Figure 7

	(1)	(2)	(3)	(4)
	# Kids	# Kids	Cohort Size	PA
Unemployment Rate				0.027
				0.007
# Kids				0.004
				0.007
Cohort Size				-0.033
				0.009
Age 30-32	omitted	omitted		
Age 33-36	0.052	-0.028		
	0.143	*0.157*		
Age 37-40	0.432	0.386		
	0.140	*0.160*		
Age 41-44	0.260	0.279		
	0.142	*0.161*		
Age 45-48	-0.343	-0.280		
	0.144	*0.170*		
Age 49-52	-0.633	-0.487		
	0.143	*0.171*		
Age 53-56	-1.034	-0.775		
	0.146	*0.179*		
Age 57-60	-1.219	-0.837		
	0.152	*0.195*		
Cohort 1896-1899		omitted	omitted	
Cohort 1900-1903		0.258	0.396	
		0.322	*0.182*	
Cohort 1904-1907		0.413	1.053	
		0.302	*0.169*	
Cohort 1908-1911		0.397	1.730	
		0.302	*0.166*	
Cohort 1912-1915		0.539	2.289	
		0.311	*0.166*	
Cohort 1916-1919		0.728	2.437	
		0.301	*0.163*	
Cohort 1920-1923		0.883	2.861	
		0.299	*0.161*	
Cohort 1924-1927		0.994	2.617	
		0.306	*0.162*	
Cohort 1928-1931		1.309	2.031	
		0.314	*0.166*	
Cohort 1932-1935		1.034	1.324	
		0.323	*0.170*	
Cohort 1936-1939		0.640	1.121	
		0.335	*0.177*	
Cohort 1940-1943		0.745	2.076	
		0.341	*0.179*	
Cohort 1944-1947		0.404	3.573	
		0.361	*0.189*	
Cohort 1948-1951		0.452	4.860	
		0.418	*0.218*	
Constant	1.904	1.079	6.490	0.508
	0.110	*0.324*	*0.153*	*0.088*
R2	0.122	0.149	0.475	0.017

notes:
standard errors in italics

Table 4: Misspecification F-Tests for Models in Figures 7-9

	Reference	Description	df	Test Statistic	p-value
		Figure 7 (Table 3)			
(1)	**Model 4**	*Basic for PA:* Compare to model with year and age terms	11, 2026	7.55	< 0.001
(2)	**Model 4**	*Full for PA:* Compare to model with year, age, and year*age interactions	45, 1992	2.69	< 0.001
(3)	**Model 4**	*Interactions for PA:* Compare model with year and age to model with age, year, and year*age interactions	34, 1992	1.11	0.304
(4)	**Models 1 & 2**	*Presence of Effect of Cohort on # Kids (arrow a3):* Compare model 1 to model 2	13, 2020	5	< 0.001
		Figure 8 (Table 5)			
(5)	**Model 3**	*Basic for PA:* Compare to model with year and age terms	12, 2026	8.13	< 0.001
(6)	**Model 3**	*Full for PA:* Compare to model with year, age, and year*age interactions	47, 1991	2.78	< 0.001
(7)	**Model 3**	*Interactions for PA:* Compare model with year and age to model with age, year, and year*age interactions	35, 1991	0.94	0.569
		Figure 9 (Table 6)			
(8)	**Model 7**	*Basic for PA:* Compare to model with year and age terms	10, 2024	1.47	0.144
(9)	**Model 7**	*Full for PA:* Compare to model with year, age, and year*age interactions	45, 1989	1.01	0.463
(10)	**Model 7**	*Interactions for PA:* Compare model with year and age to model with age, year, and year*age interactions	35, 1989	0.87	0.681
(11)	**Model 4**	*Basic for Church Attendance:* Compare to model with age terms	12, 2024	4.99	< 0.001
(12)	**Model 4**	*Full for Church Attendance:* Compare to model with age and year*age interactions	47, 1989	1.91	< 0.001
(13)	**Model 5**	*Full for Education: Compare to Model with age*cohort interactions*	27, 1993	0.77	0.794
(14)	**Model 2**	*Basic for Employed:* Compare to model with year and age terms	11, 2026	3.18	< 0.001
(15)	**Model 2**	*Full for Employed* Compare to model with year, age, and year*age interactions	45, 1992	1.55	0.012
(16)	**Model 3**	*Basic for # Kids:* Compare to model with age terms	13, 2019	4.71	< 0.001
(17)	**Model 3**	*Full for # Kids:* Compare to model with age and year*age interactions	40, 1993	8.32	< 0.001
(18)	**Model 6**	*Basic for Marital Status:* Compare to model with cohort terms (Married vs. Single)	13, 1891	1.21	0.263
(19)	**Model 6**	*Basic for Marital Status:* Compare to model with cohort terms (Married vs. Divorced)	13, 1885	3.46	< 0.001
(20)	**Model 6**	*Basic for Marital Status:* Compare to model with cohort terms (Married vs. Widowed)	13, 1831	1.01	0.436
(21)	**Model 6**	*Full for Marital Status:* Compare to model with coohrt and cohort*age interactions (Married vs. Single)	40, 1865	0.98	0.511
(22)	**Model 6**	*Full for Marital Status:* Compare to model with cohort and cohort*age interactions (Married vs. Divorced)	40, 1859	2.14	< 0.001
(23)	**Model 6**	*Full for Marital Status:* Compare to model with cohort and cohort*age interactions (Married vs. Widowed)	40, 1805	0.88	0.6901

Table 5 : OLS Estimates for Models in Figure 8

	(1)	(2)	(3)
	Employed	Church Attendance	PA
Employed			-0.162
			0.044
Church Attendance			-0.055
			0.009
Period 1956	omitted		
Period 1960	-0.028		
	0.017		
Period 1964	-0.009		
	0.016		
Period 1968	-0.019		
	0.017		
Period 1976	-0.080		
	0.017		
Period 1980	-0.046		
	0.018		
Age 30-32	omitted	omitted	
Age 33-36	-0.006	-0.004	
	0.022	*0.114*	
Age 37-40	-0.012	0.028	
	0.022	*0.116*	
Age 41-44	0.005	0.006	
	0.022	*0.117*	
Age 45-48	-0.040	-0.365	
	0.022	*0.123*	
Age 49-52	-0.032	-0.280	
	0.022	*0.124*	
Age 53-56	-0.077	-0.288	
	0.023	*0.130*	
Age 57-60	-0.122	-0.398	
	0.024	*0.141*	
Cohort 1896-1899		omitted	
Cohort 1900-1903		-0.530	
		0.233	
Cohort 1904-1907		-0.246	
		0.219	
Cohort 1908-1911		-0.268	
		0.219	
Cohort 1912-1915		-0.261	
		0.226	
Cohort 1916-1919		-0.341	
		0.218	
Cohort 1920-1923		-0.463	
		0.217	
Cohort 1924-1927		-0.562	
		0.222	
Cohort 1928-1931		-0.559	
		0.228	
Cohort 1932-1935		-0.765	
		0.234	
Cohort 1936-1939		-0.867	
		0.243	
Cohort 1940-1943		-0.931	
		0.247	
Cohort 1944-1947		-1.025	
		0.262	
Cohort 1948-1951		-1.304	
		0.303	
Constant	1.001	3.398	0.635
	0.019	*0.235*	*0.048*
R2	0.041	0.037	0.025

notes:
standard errors in italics

Table 6: OLS Estimates for Models in Figure 9

	(1)	(2)	(3)	(4)	(5)	(6)			(7)
	Cohort Size	Employed	# Kids	Church Attendance	Education	Marital Status (vs. Married)			PA
						Single	Divorced	Widowed	
Employed									-0.090
									0.042
# Kids				0.066					0.001
				0.016					*0.006*
Church Attendance									-0.033
									0.009
Education		0.008							-0.048
		0.002							*0.004*
Republican Pres.									-0.134
									0.023
Watergate									0.163
									0.026
Unemployment Rate		-0.013							
		0.004							
Cohort Size		0.023	-0.086		-0.505	0.000	-0.041	0.011	
		0.004	*0.044*		*0.363*	*0.007*	*0.013*	*0.013*	
Marital Status: Married				omitted					
Marital Status: Single				-0.195					
				0.115					
Marital Status: Divorced				-0.430					
				0.117					
Marital Status: Widowed				0.299					
				0.179					
Period 1956	omitted			omitted					
Period 1960	-0.365			0.049					
	0.010			*0.079*					
Period 1964	-0.853			-0.100					
	0.010			*0.074*					
Period 1968	-1.469			-0.279					
	0.011			*0.080*					
Period 1976	-2.768			-0.386					
	0.012			*0.080*					
Period 1980	-3.420			-0.341					
	0.013			*0.084*					
Age 30-32			omitted		omitted	omitted	omitted	omitted	
Age 33-36			0.015		-0.153	-0.005	-0.015	0.017	
			0.144		*0.280*	*0.022*	*0.043*	*0.043*	
Age 37-40			0.363		-0.500	-0.040	-0.088	0.040	
			0.144		*0.389*	*0.021*	*0.042*	*0.042*	
Age 41-44			0.168		-0.707	-0.058	-0.067	0.046	
			0.149		*0.543*	*0.022*	*0.044*	*0.044*	
Age 45-48			-0.460		-0.923	-0.050	-0.053	0.065	
			0.157		*0.751*	*0.023*	*0.046*	*0.046*	
Age 49-52			-0.799		-1.033	-0.056	-0.150	0.115	
			0.167		*0.986*	*0.025*	*0.049*	*0.049*	
Age 53-56			-1.248		-1.634	-0.056	-0.149	0.197	
			0.183		*1.209*	*0.028*	*0.054*	*0.054*	
Age 57-60			-1.496		-2.010	-0.049	-0.094	0.118	
			0.209		*1.474*	*0.031*	*0.061*	*0.062*	
Cohort 1896-1899	omitted				omitted				
Cohort 1900-1903	0.537				-0.108				
	0.028				*0.467*				
Cohort 1904-1907	1.405				-0.081				
	0.026				*0.448*				
Cohort 1908-1911	2.280				0.777				
	0.025				*0.477*				
Cohort 1912-1915	2.893				0.772				
	0.026				*0.487*				
Cohort 1916-1919	3.330				1.004				
	0.025				*0.457*				
Cohort 1920-1923	4.015				1.532				
	0.025				*0.472*				
Cohort 1924-1927	4.053				1.487				
	0.025				*0.444*				
Cohort 1928-1931	3.813				1.098				
	0.026				*0.535*				

Cohort 1932-1935	3.376				1.147				
	0.027				*0.801*				
Cohort 1936-1939	3.849				1.002				
	0.028				*0.877*				
Cohort 1940-1943	5.147				1.742				
	0.029				*0.676*				
Cohort 1944-1947	6.688				2.298				
	0.030				*0.527*				
Cohort 1948-1951	8.280				2.734				
	0.035				*0.683*				
Constant	6.490	0.694	2.762	2.811	16.589	1.095	1.525	0.888	1.104
	0.023	*0.049*	*0.457*	*0.056*	*3.830*	*0.069*	*0.133*	*0.135*	*0.069*
R2	0.988	0.029	0.123	0.045	0.104	0.009	0.011	0.015	0.114

notes:
standard errors in italics

Table 7: Example Calculation of Total Age, Period and Cohort Effects for Figure 9

Compare those born in 1936-1939 and surveyed in 1976 (49 cases) with those born in 1908-1911 and surveyed in 1960 (41 cases)

<u>Total Period Effect (1976 vs. 1960)</u>		
Watergate	1 * 0.163	0.16300000
Republican President	-1 * -0.134	0.13400000
Unemployment Rate*Employed	0.9 * -0.013 * -0.090	0.00105300
Cohort Size*Employed	-2.40 * 0.023 * -0.090	0.00496800
Cohort Size*Education	-2.40 * -0.505 * -0.048	-0.05817600
Cohort Size*# Kids	-2.40 * -0.086 * 0.001	0.00020640
Cohort Size*# Kids*Church Attendance	-2.40 * -0.086 * 0.066 * 0.006	0.00008173
Cohort Size*MS(divorced vs. married)*Church Attendance	-2.40 * -0.041 * -0.430 * -0.033	0.00139630
Cohort Size*MS(widowed vs. married)*Church Attendance	-2.40 * 0.011 * 0.299 * -0.033	0.00026049
Cohort Size*MS(single vs. married)*Church Attendance	-2.40 * 0.000 * -0.195 * -0.033	0.00000000
Church Attendance	-0.435 * -0.033	0.01435500
TOTAL		**0.2611**
<u>Total Cohort Effect (1936-1939 vs. 1908-1911)</u>		
Cohort Size*Employed	1.57 * 0.023 * -0.090	-0.00324990
Cohort Size*Education	1.57 * -0.505 * -0.048	0.03805680
Cohort Size*# Kids	1.57 * -0.086 * 0.001	-0.00013502
Cohort Size*# Kids*Church Attendance	1.57 * -0.086 * 0.066 * 0.006	-0.00005347
Cohort Size*MS(divorced vs. married)*Church Attendance	1.57 * -0.041 * -0.430 * -0.033	-0.00091341
Cohort Size*MS(widowed vs. married)*Church Attendance	1.57 * 0.011 * 0.299 * -0.033	-0.00017040
Cohort Size*MS(single vs. married)*Church Attendance	1.57 * 0.000 * -0.195 * -0.033	0.00000000
Education	0.225 * -0.048	-0.01080000
TOTAL		**0.0227**
<u>Total Age Effect (37-40 vs. 49-52)</u>		
Education	0.533 * -0.048	-0.02558400
# Kids	1.162 * 0.001	0.00116200
# Kids*Church Attendance	1.162 * 0.066 * -0.033	-0.00253084
MS(divorced vs. married)*Church Attendance	0.062 * -0.430 * -0.033	0.00087978
MS(widowed vs. married)*Church Attendance	-0.075 * 0.299 * -0.033	0.00074003
MS(single vs. married)*Church Attendance	0.016 * -0.195 * -0.033	0.00010296
TOTAL		**-0.0252**
GRAND TOTAL		**0.2586**