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

Christopher Winship

and

David J. Harding

Harvard University

(word count, not including tables and figures: 14, 177)

April 2004

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. We are grateful to Gary King, Kenneth Bollen, Robert Mare, Stephen Morgan and members of Harvard’s Applied Statistics Colloquium for comments on an earlier draft of this paper.

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

Abstract

The suggested strategy, to date, for identification of APC models has been to specify some set of restrictions on the Age, Period, and Cohort parameters that allow the model to be identified. The existing literature expresses dissatisfaction with these methods. This paper offers an alternative approach to the problem of identification that builds on Pearl's work on nonparametric causal models, in particular his front-door criterion for the identification of causal effects. The core of the estimation strategy is to specify the mechanisms through which the APC variables work. This approach allows for a broader set of identification strategies than has typically been considered in the literature and, in certain circumstances, for model specification tests. 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, 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 Age = Period – Cohort.[1]

The past literature on the identification of APC models has a number of problems. 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. 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 are sensitive to misspecification of those restrictions (Glenn 1976). Identifying restrictions are rarely if ever tested. Finally, the previous literature has not examined what it means to specify an APC model as a causal model.

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 a 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. The consequence of this is

[1] Of course, not all 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

that by adding variables to the model, identification may be possible. To put this in other terms, adding variables implies that we are adding data. With additional data, identification may be possible

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, p. 144-5)

We suggest that what is needed instead is a powerful and 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 those variables that specify the mechanisms through which Age, Period, and Cohort are assumed to work. By adding mechanisms to a model, one is harnessing the data in the mechanism variables to achieve identification.

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 being unable to estimate the parameters of interest because of model underidentification. Whereas IV involves adding variables (and thus data) 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

---

variables.

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 and the availability of measures of specific variables. Thus, as with IV, in some cases our approach will work and others it will not. However, in contrast to IV, 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 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 considerably 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 on parameter restrictions. This leads to a distinctly different and new way of thinking about identification.

Near the end of the paper we present an empirical example where the ultimate dependent variable is political alienation in order to illustrate what our approach achieves. Specifically, we use our example to illustrate how:

> 1. We have provided an explicit theoretical strategy for identifying APC models. This strategy involves specifying the mechanisms by which Age, Period, and Cohort affect the dependent variable.
>
> 2. Our approach points to a much broader set of identification strategies. Specifically, we illustrate that it is possible to have models in which:
>
> > a. More than one mechanism is associated with Age, Period, or Cohort.
> >
> > b. Age, Period, or Cohort share a mechanism.
> >
> > c. There are mechanisms that contain a component that is independent of Age, Period, and Cohort which provides a generally unrecognized and potentially powerful source of identification.
>
> 3. By considering more complicated APC models, it is possible to carry out a variety of different model identification tests. Such tests are critical in that allow the researcher to test the plausibility of the model specification.

While we focus on APC models in this paper, the general approach we use can be applied to a wide class of problems in which there are substantively distinct but linearly (or more generally functionally) dependent explanatory variables. One general class of models with identification problems are "multiple clock" problems such as the APC model. Two other examples 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. Linear dependence is also a problem for other classes of models such as status inconsistency models that attempt to assess the importance 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 father and son's status, and the mobility represented by their difference.[2]

In the next section of the paper we briefly discuss previous research. Following this, we discuss Pearl's three 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 criterion can be used to deal with unobserved variables. The subsequent section examines various specification tests. We then present our empirical example.

Limitations of Previous Research

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 three ways. First, identification can be achieved if only two of the three APC variables are assumed to affect the outcome. 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 (e.g. Firebaugh et al. 1988, Glenn 1994, Meyers et al. 1998). This is a very strong theoretical assumption which may or may not be justified in particular circumstances.

Second, as suggested in 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 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

[2] We are indebted to Robert Mare for this observation.

by Nakamura (1986) who uses a Bayesian approach to specifying restrictions. (See Sasaki and Suzuki 1987 for an application).

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 (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 through 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.

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, restrictions are typically never tested. Generally, this is because the models considered are just identified.

A fourth issue which has not been considered in the literature as far as we are aware is whether it makes sense to talk about APC models as causal models. From a counterfactual perspective (for reviews, see Winship and Morgan 1999, Winship and Sobel 2004), the assertion that Age, Period, or Cohort have causal effects is highly problematic. Holland (1986), as well as others, has argued that only manipulable variables can have a causal effect (for further discussion see Winship and Sobel 2004). 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 be changed exogenously. Similarly, there is no way to exogenously change the relevant period.

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 here is that a cause must have the ability to bring about an effect. (Bunge 1979, Harrẹ 1972, Harrẹ and Madden 1975). This is posited to occur because the cause is related to the outcome through some set of mechanisms (Cartwright 1989, Glennan 1996). For example, although there is no counterfactual experiment that can be directly done to show that the moon causes tides, we believe that there is a causal relationship between the two due to the asserted mechanism involving the gravitational pull of the moon. Conversely, although night always follows the day, we do not believe that day causes the night because of the absence of a causal mechanism by which day might cause night, and our belief that both are related to a separate causal factor, the rotation of the earth. (For a more detailed discussion of the different concepts of causality, see Sobel 1995, Brady 2003, and Winship and Sobel 2004.)

Assuming we can specify the mechanisms through which Age, Period, or Cohort affect an outcome, does it make sense to talk about the mechanisms involved as having causal effects? Often, if not always, the mechanisms through which Age, Period, and Cohort are assumed to work are manipulable. For example, we might be interested in the importance of aging because of its association with 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 different 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.

Are APC models then causal models? We would argue that since Age, Period, and Cohort clearly are not manipulable, it only makes sense to think of APC models as causal models if we can conceptually specify the mechanism through which Age, Period, and Cohort work. If this can be done and the associated mechanisms are manipulable, then APC models can, in a well defined sense, can be thought of as causal models. Thus the specification of the mechanisms through which Age, Period, and Cohort work is not only critical for identification, but also for creating a well-defined causal model.

The 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. 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: in order to achieve identification of APC models and have a causally well-defined model, it is necessary to specify the mechanisms through which the processes of interest work.[3]

[3] Methodologists working on APC models have always advocated the use of theory in the identification of such models, but previously theory was primarily used to justify the exclusion of either Age, Period, or Cohort from the model, setting two or more coefficients to be equal, or the use of a particular proxy variable. Our approach departs

The essential point in both the theoretical literature and this discussion are the same: in order to know why two events are associated, 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 *nonparametric* 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 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.

In many ways, Pearl's theory is similar to the standard theory of linear paths 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). 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. Pearl's approach involves a calculus of conditional probabilities that is analogous to, but generalizes the standard path calculus. Since an understanding of this conditional probability calculus is not essential to this paper, we do not develop it here. The reader is referred to Pearl's 2000 book for details (also see Pearl 1999).

The general problem that Pearl is concerned with is distinguishing true causation from

---

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.

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. In the figures we discuss below, however, we omit the error terms. Although our approach is applicable to a wide range of models, in all of the hypothetical examples we consider as well as in our empirical example, we only consider 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.

Figure 1 shows an extremely simple example where X and Y are directly connected by a "backdoor" path through Z_1 and Z_2. 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. Here the goal is to estimate the direct effects of a particular independent variable on some outcome. 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 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 simply estimated by conditioning on one of the variables in each path. 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 the *Z*'s. Deleting ***either*** Z_1 ***or*** Z_2 from this graph, which in Pearl's theory 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*. This conditioning strategy is only possible if the *Z*'s 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 is 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.[4] 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.

[4] 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).

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 total effect of *S* on *C*. The covariance/correlation between S and C 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 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* 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 back-door paths between *S* and *T*, we can consistently estimate the effect of *S* on *T*. There is, however, a backdoor path between *T* and *C* through *S* and *U*. However, by conditioning on *S*, we can eliminate this backdoor path, which allows us to consistently estimate the effect of *T* on *C*.

The example that Pearl (2000) gives is of the effects of smoking (*S*) on cancer (*C*). 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 increases the risk of cancer primarily through the presence of tar, then we can estimate the effect of smoking on tar, *b*, and the effect of tar on cancer, *c* to get an estimate of the effect of smoking on cancer, the product of the two effects – *bc*.

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 *T* 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 cancer (*C*).

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. We argue, here, however, that the front door criterion 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.

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. By adding variables we are obviously adding additional data to the analysis. In most cases 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, 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 there is 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. Within the current context, this would mean that the independent variables within each equation must be linearly independent. Those familiar with structural equation models will see parallels between that framework and the mechanics of what we propose here, but in order to keep the presentation accessible to the widest audience, we chose not to rely on the language of SEM.

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

Y = Outcome of Interest

X is a matrix consisting of a constant plus the variables:

Age = *A* = Years Since Date of Birth

Period = *P* = Current Year

Cohort = *C* = Year of Birth

Error = *e* = error term in the APC regression equation.

a is a 3 by 1 vector of parameters to be estimated, where each parameter measures the relative contribution of Age, Period, and Cohort to change in *Y*.

Our goal is to estimate:

$$(1)\ \ Y = \mathbf{Xa} + e = a_0 + \mathbf{A}a_1 + \mathbf{P}a_2 + \mathbf{C}a_3 + e,$$

but, because of the identity Age = Period – Cohort, $(X'X)^{-1}$ does not exist and as a result equation (1) cannot be estimated by Ordinary Least Squares. The same issue would exist if Age, Period, and Cohort were specified in terms of a set of dummy variables.

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 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. 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 straightforwardly estimate the parameters associated with Age, Period, and Cohort because of their linear dependence.

Now generalize our model by assuming that there is a matrix of *m* variables, **M**, that represent the mechanisms associated with Age, Period, and Cohort that affect the outcome *Y*. Also let **B** be an *m* by 3 matrix of parameters that specify the relationship between Age, Period, and Cohort. Let **U** be an *n* by *m* matrix of errors in the equation specifying the relationships between **M** and **X**. Let **c** be a 1 by *m* vector of parameters to be estimated that represent the effect of each mechanism on the outcome **Y**. We then have the following set of equations:

(1) $Y = \mathbf{Xa} + e$	Equation specifying the relationship between the Outcome (*Y*) and Age, Period, and Cohort (**X**)
(2) $\mathbf{M} = \mathbf{XB} + \mathbf{U}$	Set of m equations specifying the relationship among the mechanisms (**M**) and Age, Period, and Cohort (**X**)
(3) $Y = \mathbf{Mc} + v$	Equation specifying the relationship between the outcome (*Y*) and the mechanisms (**M**)

Substituting (2) into (3), we get:

(4) $Y = \mathbf{XBc} + \mathbf{Uc} + v$	Reduce form equation specifying the relationship between the outcome (*Y*) and Age, Period, and Cohort (**X**)

If we can estimate **B** and **c**, then we can estimate **a** = **Bc**.

In the above equations two conditions are necessary, though they may not be sufficient, for identification of the overall model. First, any mechanism in **M** that is directly affected by Age, Period, and Cohort must either be affected by at most two of these three variables or there must be some parameter restriction that will allow the identification of the equation for which this mechanism variable is the dependent variable. If neither of these conditions holds, identification may be possible by treating the mechanism as an outcome in its own right and attempting to find other intermediary variables, i.e. mechanisms that affect this mechanism.

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. If neither of these conditions hold, then it may be possible to achieve identification by adding intermediate variables, that is, additional mechanisms, between the mechanisms and the outcome. 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. Although we do not formally investigate models where there are mechanisms that affect each other, we do present empirical examples where this is the case.

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- characteristic model. In our model this is equivalent to specifying the mechanisms that are associated respectively with Age, Period, and Cohort, that affect the outcome. As an example, the unemployment rate might be the mechanism associated with period. Cohort size might be the mechanism associated with birth cohort. As long as these new variables are not linearly dependent, then it will be possible to estimate their effects on the outcome. Generally, this will be the case if these variables are not strict linear functions of Age, Period, and Cohort. Putting it differently, 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 M_A be a variable representing the mechanism associated with Age, and let M_P and M_C 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 use proxy variables to represent Age, Period, and Cohort is that this is an application of Pearl's front door criterion. In the model in Figure 4, in general we should be able to estimate **b** coefficients since there is no problem of linear dependence or, if the relationship is deterministic, specify them. A necessary condition for estimating the **c** coefficients, that is, the effects of M_A, M_P, and M_C, is that they 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*, using the appropriate products. We need, however, to consider the necessary conditions for the *M* 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 separate sets of dummy variables, D_A, D_P, and D_C, for every separate value of *A*, *P*, and *C*. and analogously redefine M_A, M_P, and M_C as sets of dummy variables. If the separate values of the **D** variables map onto unique values of the **M** 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 **M** variables that exist 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 **M** variables and the outcome *Y*, for example as linear, then identification can be achieved as long as the **M**'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 **M** variable. In this case, the specification of a relationship between a **D** variable and an **M** 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 **M** 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 in order 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 of 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 APC 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 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 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 *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*.

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 might 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, 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 we have identified all the mechanisms for ***one*** of the APC variables. 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 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 criteria states that if we can consistently estimate the **b** and **c** 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 **b** coefficients is unproblematic. Also assume that the **M** functions are ***not*** 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 UM_A and UM_C) connecting each **M** variable and *Y*.

Consider the problem of estimating the effect of M_P on *Y*, c_2. There are 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, 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 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* 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 b_1 conditioning on *A*. There will be an interdependence problem. Let's say, however, that there is variation in M_A independent of *A*. This would generally be true for a variable like marital status. 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 down 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 in order 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 marital status example 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 variables on which it depends, then its effect can also be identified by conditioning on those variables, as long as the intermediary is not an exact deterministic function of the variables on which it depends.

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 dependence problem between these three variables. We could, however, break these backdoor paths by conditioning on M_P and C. <u>Note that this demonstrates that 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</u>. Thus, in this example, it is possible to identify all three effects if all the effects associated with P are observed.

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 argue that specification testing should be done in two stages. The first goal should be to test the overall adequacy of the model. Two related tests are available for doing this. Both tests assess whether the mechanisms that have been posited fully explain the relationship between Age, Period and Cohort and the outcome. In the second stage, the identifying assumption(s) for each mechanism should be tested.

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 M_A, M_P, and M_C and calculate R-square.

Step 2: Regress Y on M_A, M_P, M_C, D_A, and D_P and calculate R-square.

Step 3: Perform an F-test.

If we have correctly specified the mechanisms through which Age, Period, and Cohort are associated with Y, then to within sampling error, the two R-squares should be the same. This can be formally tested using an F-test since the first model is nested within the second. 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. We refer to this type of test as a *global unsaturated test*.

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. We refer to this type of test as a *global saturated test*.

The global tests can be performed whenever the model specifies mechanisms for all or all but one of APC. If we specify mechanisms for only one of APC, then there will be no larger equation to estimate in step 2 because of the linear dependence of APC. To see an example of the global unsaturated test when we have specified a mechanism for all but one of APC, consider the model in Figure 4 but imagine that we have left out mechanism M_A (replacing the pathway D_A-M_A -Y with a single arrow D_A-Y). In this case, step 1 would be to regress Y on D_A, M_P, and M_C, and step 2 would be to regress Y on D_A, M_P, M_C, and D_P. Step 3 is the same as above. Since the model in step 1 is nested within the model in step 2, we can again use an F-test to compare them, examining whether there is any further variance in the outcome associated with Period or Cohort once we have entered our mechanisms in the model.

Both the global unsaturated and saturated tests should be understood as general specification tests. Specifically, they can be thought of as Hausman (1978) type tests. 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.

The above two tests assess whether we have specified a set of mechanisms that fully explains the relationship between Age, Period, and Cohort and the outcome (or two of the three if we do not include a mechanism for one of the three). In general, this should test whether we have specified the full set of mechanisms through which they work. The exception to this statement would be if there were a missing mechanism that was a deterministic function of the other mechanisms in the model. This situation should be unusual.

What the above tests do not do is assess whether we have appropriately specified the particular relationship between Age, Period, and Cohort and each of the mechanisms. As discussed above, these relationships will only be identified if each mechanism is affected by at most two of APC. These assumptions cannot be tested, but must be justified by theory. The preferable situation is when each mechanism is assumed to be uniquely associated with Age, Period, or Cohort. This assumption can be directly tested. In this case, we can use standard F tests or Chi-square likelihood ratio tests to compare the models. For example, in Figure 4 above, we can test whether M_A is affected by Period as well as Age using either a t-test or an F-test. We refer to such a test as a ***mechanism test***.

This discussion demonstrates that, in a model where each mechanism is assumed to be affected by only one of Age, Period, or Cohort, all the relationships in the model can be empirically tested. Figure 4 is an example of such a model. As discussed above, we need to start by carrying out a general specification test to determine whether the posited mechanisms do in fact capture all of the effects of Age, Period, and Cohort on the outcome. If this hypothesis cannot be rejected, then we need to individually test whether each mechanism is affected by only one of Age, Period, and Cohort. If this set of hypotheses cannot be rejected, then we can assert that the model has been fully empirically tested.[5] On the other hand, if the model specifies that a mechanism is affected by two of APC, then this portion of the model cannot be tested, and we must rely on theory for its justification. Figure 5 is an example of such a model.

Empirical Example

To illustrate these ideas, we conduct a basic analysis of the effects of Age, Period, and

[5] Note that this situation is distinctly different from instrumental variables, in which it is never possible to determine whether a single instrument is valid, but only whether a set of two or more provide consistent estimates of the same value.

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.[6] 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 surveyed in 1956, 1960, 1964, 1968, 1976 and 1980 who have no missing data on any of our variables, leaving 2041 cases.[7] Throughout this example, we use linear probability models for maximum accessibility and expository clarity. Such models may not be appropriate for other applications, but here results are similar when probit, logit, ordered probit, or multinomial logit models are used instead of OLS.

--- Table 1 here ---

--- Figure 7 here ---

Figure 7 shows a simple model of the relationships between Age, Period, Cohort, and Political Alienation (hereafter, PA). As before, since 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 variable of interest: Age, Period, and Cohort. (For now, ignore the dashed arrow labeled b_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

[6] Data are from the National Election Studies Cumulative Data File, 1948-2000, ICPSR No. 8475 (Sapiro, Rosenstone, and the National Election Studies 2002).
[7] 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.

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 c_1, c_2, and $c_{3,}$ estimate a model predicting number of kids with Age to estimate b_2, and estimate a model predicting cohort size with cohort to estimate b_4.[8] Because the relationship between Period and Unemployment Rate is deterministic, we already know b_1 (see Table 2). The effects of Period, Age, and Cohort on PA are then (b_1c_1), (b_2c_2), and (b_4c_3), respectively.

--- Table 2 here ---

Table 2 indicates the relationship between each level of Period and Cohort and their deterministic mechanism 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 has a highly nonlinear relationship with Period.

--- Table 3 here ---

Table 3 reports parameter estimates for the model in Figure 7. 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

[8] 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.

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 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 ---

To conduct the global unsaturated test, 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 us the global saturated test. As shown in row 2 of Table 4, this test produces 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 unsaturated and saturated 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).[9] These tests suggest 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

[9] Were we using a a maximum likelihood Probit model rather than a linear probability model for PA, we could simply substitute likelihood ratio tests for the F tests.

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.

We should also use the mechanism tests to examine whether our assumption that each mechanism is only affected by one of APC is correct. For example, we might wonder whether Number of Kids is truly only a function of Age, or whether Cohort also affects Number of Kids. This additional relationship is represented in Figure 7 by the dashed arrow labeled b_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). This test shows us that our assumption that Number of Kids is only affected by Age is incorrect. Normally, we would also conduct mechanism tests for Unemployment Rate and Cohort Size, but here Unemployment Rate is a deterministic function of Period and Cohort Size is a deterministic function of Cohort, so we have constructed them in such a way that they necessarily pass such tests.

--- Figure 8 here ---

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 will 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 APC variables. The effect of A on PA operates through both employed and church attendance and is estimated as $(b_2c_1)+(b_3c_2)$. The effect of P on PA operates only through employed and is estimated as b_1c_1. And the effect of C on PA operates only through church attendance and is estimated as b_4c_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 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 global unsaturated and global saturated specification tests. 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 global unsaturated test), 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 (the global saturated test), 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. Note that we cannot perform the mechanism tests for the model in Figure 5 because each mechanism is affected by two of APC.

--- 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 global specification tests to find out. We begin with the global unsaturated test. 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 global unsaturated 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 mechanism variables between A, P, or C and PA that are not represented in our model. We also performed the global saturated test by including all interactions in the second model. 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).

Having finally come up with a model that passes the global unsaturated and saturated tests, we can now turn to the mechanism tests. These tests of intervening variables are critical. 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.

We begin with the mechanism test for Church Attendance. 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. The mechanism test examines the validity of the latter two assumptions. In conducting the test, we estimate two models and compare them. The first model predicts Church Attendance with Cohort Size, Marital Status and dummies for each Period (see Column 4 of Table 6) and embodies our assumptions. 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).

At this point it is worth pausing to consider the relationship between the global tests and the mechanism test. Though for pedagogical purposes, and to emphasize the role of mechanisms, we have distinguished the global tests from the mechanism test, statistically and conceptually they are the same. They merely differ in the variable on which they focus. The global tests focus on whether we have accounted for all mechanisms between A, P, and C and the dependent variable (here, political alienation), while the mechanism test focuses on whether we have accounted for the relationship between A, P, and C and a particular mechanism. In essence, both types of test assess whether variables not thought to directly affect the focal variable add any explanatory power to the model. Thus, we can also conduct both unsaturated and saturated versions of the mechanism test. Above, we conducted the unsaturated mechanism test for Church Attendance. Row 12 of Table 4 shows the corresponding saturated mechanism test by including

Age-Period interactions in the second model as well. This test too is significant (F = 1.91 with 47 and 1989 degrees of freedom, p < 0.001). The results of these two tests suggest that we have not captured all the mechanisms connecting A, P, and C and Church Attendance with our model.

We can do the unsaturated mechanism 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 except for the model containing interactions between the two. For example, in Figure 9 Education is directly affected by both Cohort and Age, so we can only conduct a weak version of the saturated mechanism test. 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). This result tells us that the Age by Cohort interaction is not necessary, but it cannot tell us whether our assumption that period does not directly affect education is correct or not. There is no way to test this assumption.

We now turn to mechanism tests for the remaining mechanisms. Unfortunately, as shown in Table 4, Employed (rows 14 and 15) and Number of Kids (rows 16 and 17) both fail both their mechanism 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 mechanism tests, though the other Marital Status comparisons pass (see rows 18 to 23 in Table 4).

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.

The results of our mechanism tests mean we have failed to specify at least one of the paths between A, P, or C and each of the intervening variables. How might we go about improving the model in Figure 9? We could add more mechanism variables to further specify the relationships between A, P, and C and the mechanism variables we already have in the model. This is the ideal solution because, by building up the model, we move toward a model in which each mechanism is affected by at most one of APC. This model would be ideal because it would be fully testable. Unfortunately, there are no more variables in the data to help us to fill out the relationships between APC and the mechanisms already in the model. 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 APC variables. A second best solution would be to add further arrows between APC and our existing mechanisms. This would fill out the model and lead to an estimable model, but it would also force us to rely heavily on the untestable assumptions that each mechanism is affected by only two of the three of APC. In addition, it would force us to rely on theoretical assumptions to decide which two of the three of APC affect each mechanism, essentially taking us back to the original problem that motivates this paper. Thus, we do pursue that solution here.

Although our model has not passed all specification tests, for illustrative purposes 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.

-- Table 7 here --

To calculate the total effect of either Age, Period, or Cohort on PA, we must first list all the paths through which each variable 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.

The use of linear probability models 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. In this paper we have presented a new methodological approach to the identification of APC models. Our method also is applicable to other "multiple clock" models where there is linear, more generally functional dependence, among variables and other models, such as models of social mobility and status inconsistency where dependence is a problem.

Our approach builds on previous work by Mason and others. We present an illustrative example to demonstrate what we have achieved. Specifically, we use our example to illustrate how:

1. We have provided an explicit theoretical strategy for identifying APC models. This strategy involves specifying the mechanisms by which Age, Period, and Cohort affect the dependent variable.

2. Our approach points to a much broader set of identification strategies. Specifically, we illustrate that it is possible to have models in which:

    a. More than one mechanism is associated with Age, Period, or Cohort.

    b. Age, Period, or Cohort share a mechanism.

    c. There are mechanisms that contain a component that is independent of Age, Period, and Cohort which provides a generally unrecognized and potentially powerful source of identification.

3. By considering more complicated APC models, it is possible to carry out a variety of different model identification tests. Such tests are critical in that allow the researcher to test the plausibility of the model specification.

The approach we consider is quite general. Our hope is that it will renew interest in APC models and other models in which linear (or functional) dependence is an issue.

References

Brady, Henry E. 2003. “Model of Causal Inference: Going Beyond the Neyman-Rubin-Holland Theory.” Unpublished. Department of Political Science, University of California – Berkeley.

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

Bunge, Mario A. 1979. *Causality and Modern Science.* (3rd ed.), New York: Dover.

Cartwright, Nancy. 1989. *Nature’s Capacities and Their Measurement.* New York: Oxford University Press.

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.

Firebaugh, Glenn and Kenneth E. Davis. 1988. "Trends in Antiblack Prejudice, 1972-1984: Region and Cohort Effects." *American Journal of Sociology* 94(2): 251-72.

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

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

Glenn, Norval D. 1994. "Television Watching, Newspaper Reading, and Cohort

Differences in Verbal Ability." *Sociology of Education* 67: 216-230.

Glennaan, Stuart S. 1996. "Mechanisms and the Nature of Causation." *Erkenntnis, V. 44*, Pp. 49-71.

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

Harrẹ, Rorn. 1972. *The Philosophies of Science.* Oxford: Oxford University Press.

Harrẹ, Rorn, and Edward H. Madden. 1975. *Causal Powers: A Theory of Natural Necessity.* Oxford: Basil Blackwell.

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.

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. 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.

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.

Sobel, Michael. 1995. “Causal Inference in the Social and Behavioral Sciences.” In *Handbook of Statistical Modeling for the Social and Behavioral Sciences.* Ed. G. Arminger, CC Clogg, and ME Sobel. Pp. 1-38. New York: Plenum.

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. 2004. “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: Backdoor Criterion

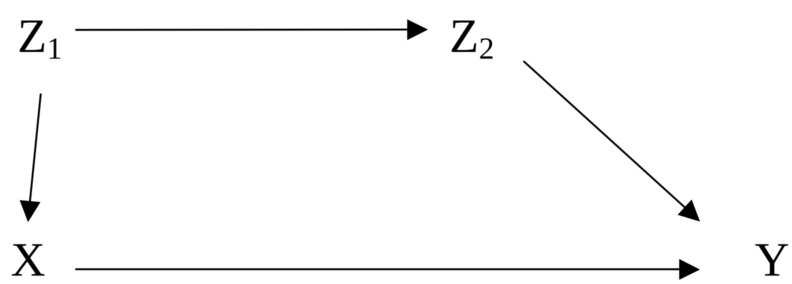

Figure 2: Instrumental Variables

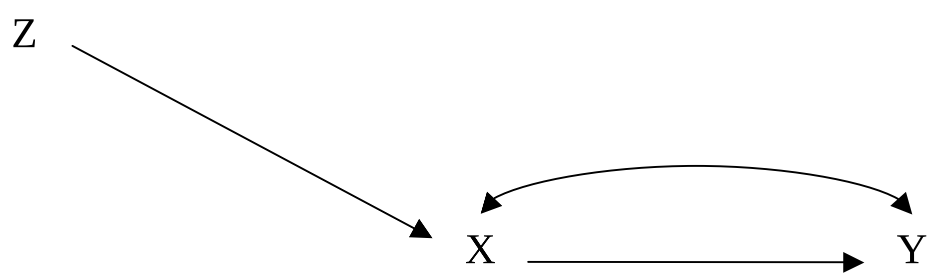

Figure 3: Front Door Criterion

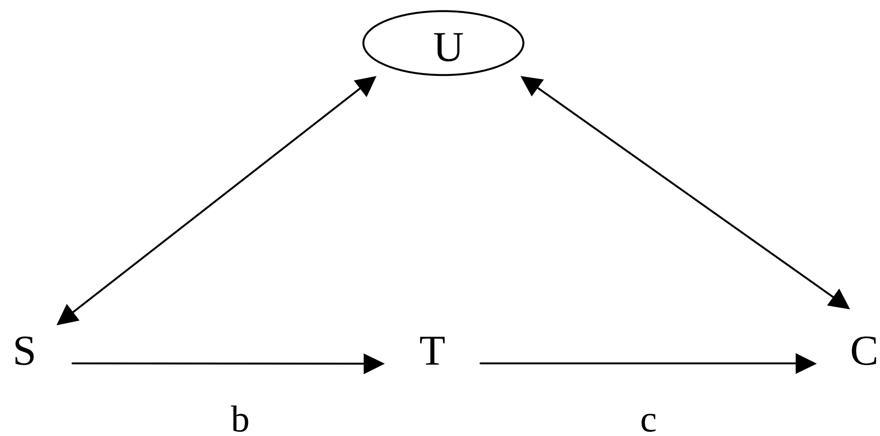

Figure 4: Hypothetical APC Model with Intervening Mechanisms

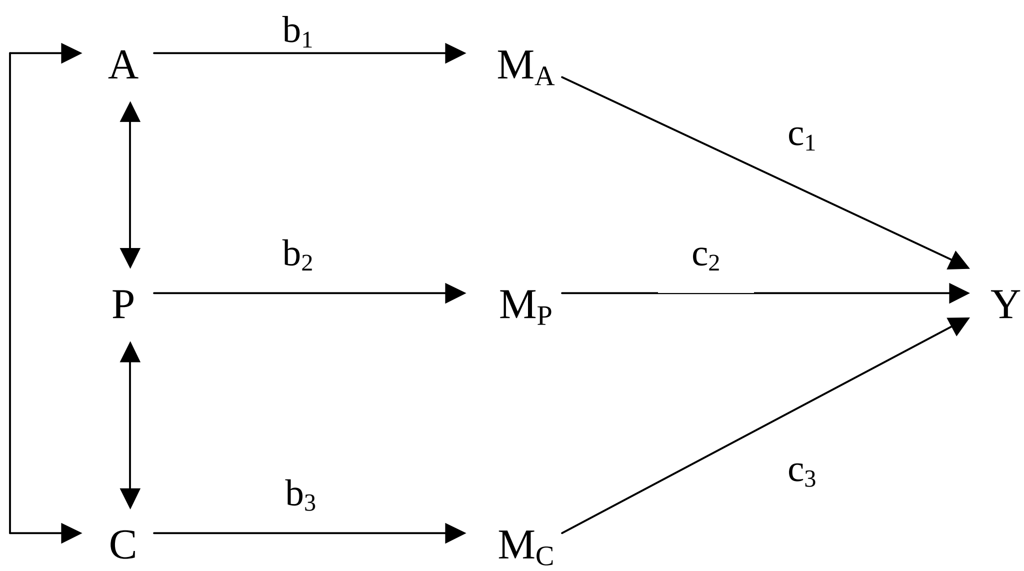

Figure 5: Hypothetical APC Model with Shared Mechanisms

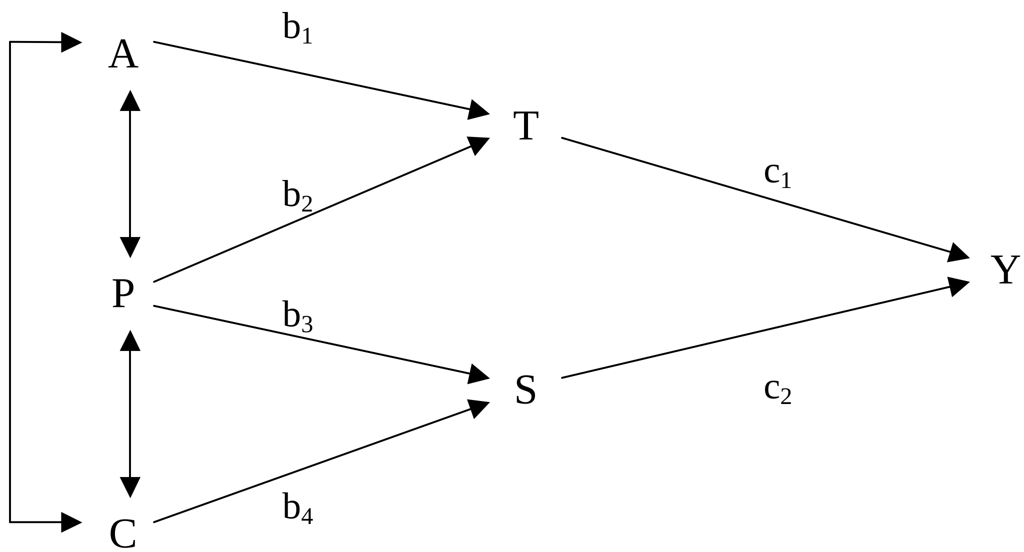

Figure 6: Hypothetical APC Model with Unobserved Mechanisms

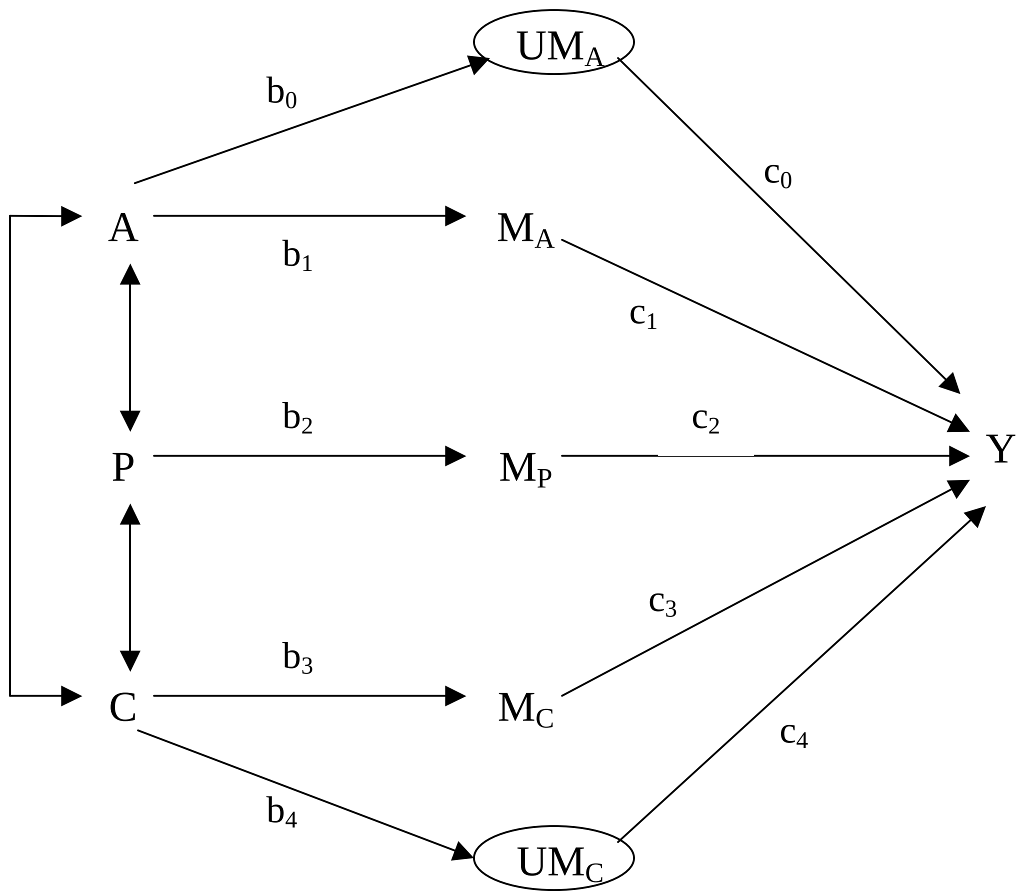

Figure 7: Simple APC Model for Political Alienation with Intervening Mechanisms

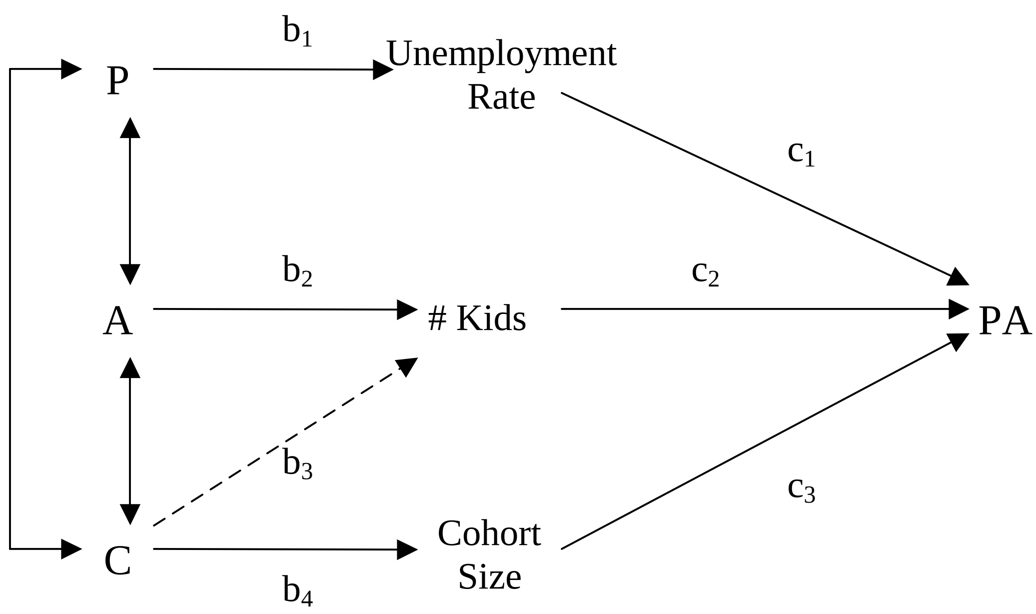

Figure 8: Simple APC Model for Political Alienation with Shared Intervening Mechanisms

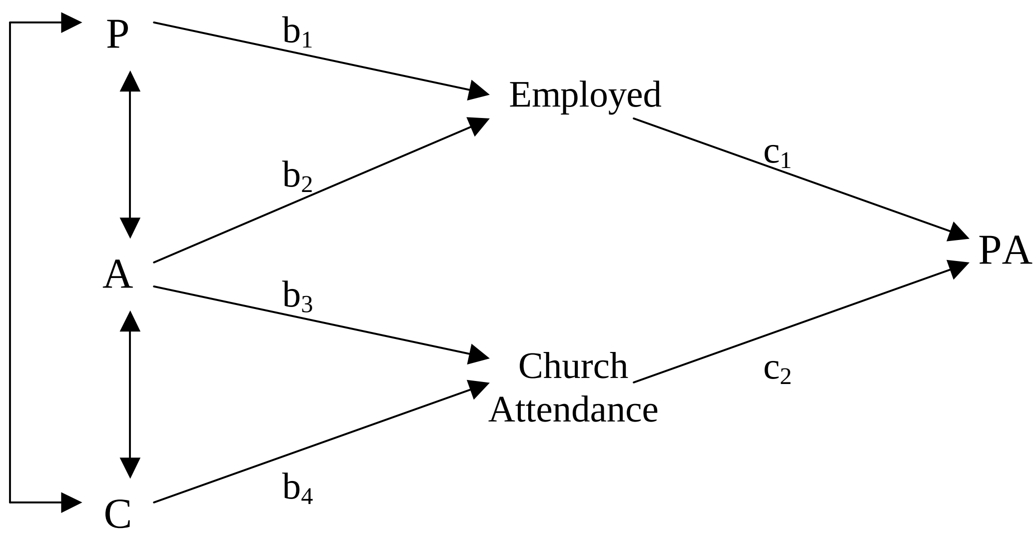

Figure 9: Full APC Model for Political Alienation with Multiple Mechanisms

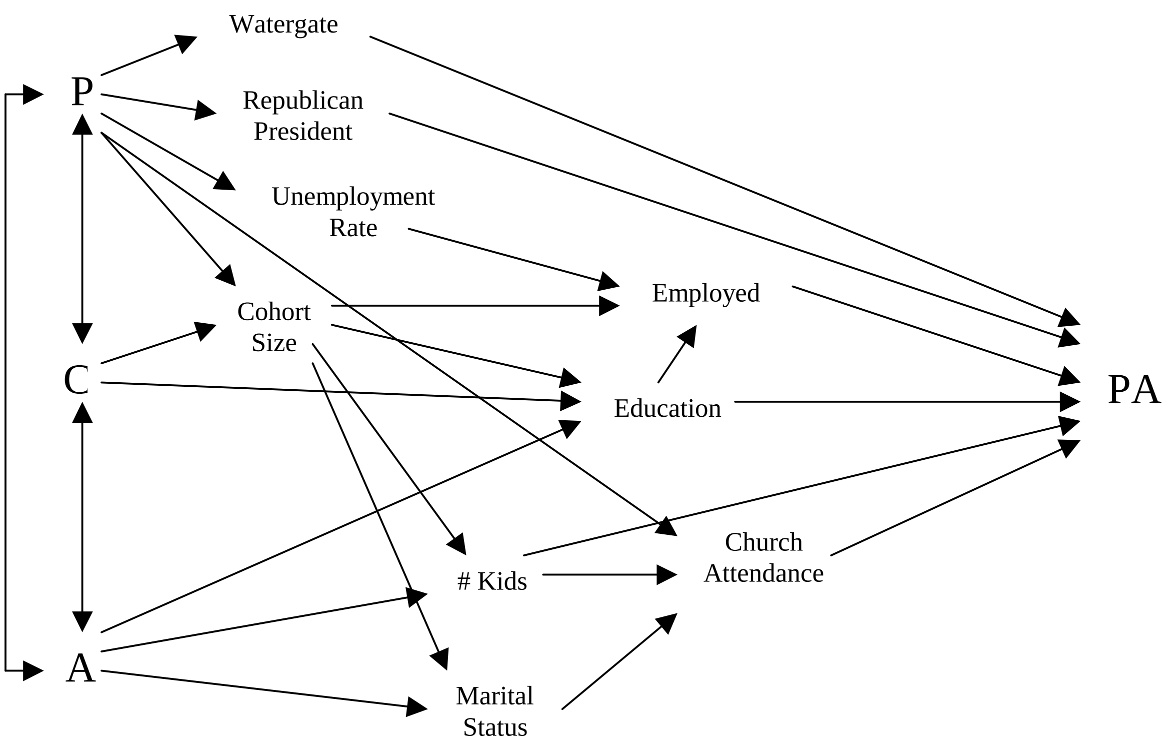

Table 1: Variable Descriptions for Political Alienation Example

Name	Description	Source
Political Alienation (PA)	Agree with statement in text.	NES (VCF0609)
Age (A)	30 to 60, in 4-year age groups	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**	*Unsaturated for PA:* Compare to model with year and age terms	11, 2026	7.55	< 0.001
(2)	**Model 4**	*Saturated 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**	*Unsaturated for PA:* Compare to model with year and age terms	12, 2026	8.13	< 0.001
(6)	**Model 3**	*Saturated 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**	*Unsaturated for PA:* Compare to model with year and age terms	10, 2024	1.47	0.144
(9)	**Model 7**	*Saturated 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**	*Unsaturated for Church Attendance:* Compare to model with age terms	12, 2024	4.99	< 0.001
(12)	**Model 4**	*Saturated for Church Attendance:* Compare to model with age and year*age interactions	47, 1989	1.91	< 0.001
(13)	**Model 5**	*Saturated for Education: Compare to Model with age*cohort interactions*	27, 1993	0.77	0.794
(14)	**Model 2**	*Unsaturated for Employed: Compare to model with year and age terms*	11, 2026	3.18	< 0.001
(15)	**Model 2**	*Saturated for Employed Compare to model with year, age, and year*age interactions*	45, 1992	1.55	0.012
(16)	**Model 3**	*Unsaturated for # Kids: Compare to model with age terms*	13, 2019	4.71	< 0.001
(17)	**Model 3**	*Saturated for # Kids: Compare to model with age and year*age interactions*	40, 1993	8.32	< 0.001
(18)	**Model 6**	*Unsaturated for Marital Status: Compare to model with cohort terms (Married vs. Single)*	13, 1891	1.21	0.263
(19)	**Model 6**	*Unsaturated for Marital Status: Compare to model with cohort terms (Married vs. Divorced)*	13, 1885	3.46	< 0.001
(20)	**Model 6**	*Unsaturated for Marital Status: Compare to model with cohort terms (Married vs. Widowed)*	13, 1831	1.01	0.436
(21)	**Model 6**	*Saturated for Marital Status: Compare to model with coohrt and cohort*age interactions (Married vs. Single)*	40, 1865	0.98	0.511
(22)	**Model 6**	*Saturated for Marital Status: Compare to model with cohort and cohort*age interactions (Married vs. Divorced)*	40, 1859	2.14	< 0.001
(23)	**Model 6**	*Saturated 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)			
						Single	Divorced	Widowed	PA
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)

Total Period Effect (1976 vs. 1960)		
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**
Total Cohort Effect (1936-1939 vs. 1908-1911)		
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**
Total Age Effect (37-40 vs. 49-52)		
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**