Causality in Life Course Studies

Ravaris Moore
Jennie Brand

PWP-CCPR-2014-018

Latest Revised: December 16, 2014

California Center for Population Research
On-Line Working Paper Series

CAUSALITY IN LIFE COURSE STUDIES*

RAVARIS MOORE
University of California – Los Angeles

JENNIE E. BRAND
University of California – Los Angeles

In Preparation for: *Handbook of the Life Course, 2nd Edition*
Michael Shanahan, Monica Johnson, and Jeylan Mortimer eds., Springer Series

Version: December 16, 2014
Running head: causality and the life course
Word count (including abstract, references, appendix): 12,712

* Direct all correspondence to Ravaris Moore, Department of Sociology, University of California – Los Angeles, 264 Haines Hall, Los Angeles, CA 90095-1551, USA; email: ravarism@ucla.edu; phone: XXX. This research made use of facilities and resources at the California Center for Population Research, UCLA, which receives core support from the National Institute of Child Health and Human Development, Grant R24HD041022. The National Institutes of Health, Grant 1 R01 HD07460301A1, provided financial support for this research. The ideas expressed herein are those of the authors.

AUTHOR BIOGRAPHY PAGE

RAVARIS MOORE is a graduate student in sociology at the University of California – Los Angeles. His work examines the effects of environmental factors on the educational and career outcomes of youth. His research beyond examines the effects of exposure to violence on the academic outcomes of youth, and the degree to which applied high school S.T.E.M. courses increase the rates at which disabled youth pursue college S.T.E.M. majors and S.T.E.M. related careers.

JENNIE E. BRAND is Associate Professor of Sociology at the University of California – Los Angeles and Associate Director of the California Center for Population Research (CCPR). Her research centers on inequality and its implications for various outcomes that indicate life chances. This substantive focus accompanies a methodological focus on causal inference and the application and innovation of statistical models for panel data. Current research projects include evaluation of heterogeneity in the effects of education on socioeconomic outcomes and the socioeconomic and social-psychological consequences of job displacement.

CAUSALITY IN LIFE COURSE STUDIES

ABSTRACT

This chapter explores common methods of causal inference used in life course research. The central premise of life course research is the presumption that no period of life can be understood in isolation from prior experiences, as well as individual's aspirations for the future. This chapter discusses the context of common causal inference methods as they relate to life course research, including regression, propensity score matching, instrumental variables, and fixed effects. We also discuss strategies for incorporating variation in response to treatment according to heterogeneity, time-variation, and mediation, important components to estimating effects over the life course with a causal framework. The chapter aims to explain the assumptions behind the methods we present, and includes some heuristic derivations to aid in intuitive explanations. We also provide examples of the methods discussed using constructed data with a known data generating process.

Keywords: causal inference; counterfactuals; life course; matching; fixed effects; instrumental variables; time-varying treatments; mediation analysis;

CAUSALITY IN LIFE COURSE STUDIES

1 Introduction

This chapter surveys common methods employed for causal inference in life course research. Pearl (2009) explains that cross-sectional data can establish associations between variables, but alone cannot establish causality. In the methods described below, causal inference is the result of estimating the conditional change in an outcome associated with changes in an independent variable in a theoretical framework where the identified relationship can be plausibly interpreted as causal. This implies that inferring causality generally rests on the researcher's ability to complete two objectives: (1) Formulate a research framework where the relationship between two measures x and y, should be interpreted as a change in x initiating a change in y, as opposed to a change in y initiating a change in x, x and y causing changes in each other simultaneously, or x and y both responding to changes in a third measure z; and (2) Generate an unbiased estimate of the change in y associated with a change in x. The primary methodological emphasis in this chapter, and indeed in the literature, is on achieving objective (2); objective (1), however, deserves discussion as well.

The most convincing frameworks for facilitating causal inference tend to argue that the independent event of interest x was governed by a process exogenous to the outcome of interest y, and that x occurred chronologically before y. The exogeneity of x guarantees the absence of an endogenous measure, z, that influences both x and y, while the chronological timing implies that y could not have caused x. This leaves a causal effect of x on y as the only plausible interpretation. Beyond experimental or quasi-experimental settings, in which these

conditions may be more convincingly satisfied, we often aim to employ conditional independence as a way of limiting potential pathways of causality. With conditional independence, we assume that after controlling for differences in some set of observables, w, x is governed by a process exogenous to y. This condition, with the continued assumption that x occurred chronologically before y can also be sufficient for causal inference. In practice, this type of argument is employed very often with the caveat that there may be additional observable or unobservable measures excluded from w that result in a violation of conditional independence, and compromised grounds for causal inference.

Below we discuss a number of methods commonly employed to understand the conditional change in an outcome associated with a change in an independent variable. This includes multivariate regression, matching models, instrumental variable models, and fixed effects models, all useful approaches for causal inference when estimating effects over the life course. We also briefly discuss experimental and quasi-experimental frameworks. To streamline the discussion, several methods were omitted such as structural equation models, growth curve models, method of moment estimators, and a large class of non-parametric estimators. For the models discussed below, we aim to give the reader a practical intuition for the usefulness and applicability of each model. In some cases we add heuristic mathematical motivation to illustrate technical points and illustrate the benefits and limitations of model assumptions. We also construct simulated data and estimate some of the listed models on our constructed data set.

The chapter proceeds as follows. Section two describes various causal inference methods and discusses their benefits and limitations in life course research. Section three extends our

basic models to incorporate response variation according to individual characteristics, time, and post-treatment events and circumstances. Section four describes our research question, simulated data, and the information available to the researcher. Section five explains the process that generated the simulated data, and discusses the true simulated relationship between treatment and outcome. Section six compares estimates from various causal inference methods in section four to the true effect estimates from section five and discusses the strengths and weaknesses of each method. Section seven concludes.

2 Methods for Causal Inference in Life Course Research

This section describes, discusses, and illustrates the use of several common causal inference approaches. We start each section with a brief statistical presentation of each model and it's necessary assumptions, and discuss some of the strengths and weaknesses of each method.

A. Regression Models

Ordinary Least Squares (OLS) regression models are undoubtedly the most heavily employed tool for understanding the relationship between vectors. OLS models the conditional expected value of a measure, y, given a fixed value of, $d.\delta$ can be estimated directly as a linear projection. Given $N\times1$ matrices y and ε, and $N\times K$ matrix X, with the assumed relationship $y = d\delta + \varepsilon$, we can solve for an estimator, $\hat{\delta}_{OLS}$:

$$\hat{\delta}_{OLS} = [d'd]^{-1}d'y$$

Cameron and Trivedi (2005) specify a set of assumptions which guarantee consistency of $\hat{\delta}_{OLS}$. Consistency roughly means that the parameter estimate $\hat{\delta}_{OLS}$ converges to δ as sample size increases. We provide a brief overview of those assumptions below.

i The data are generated by a process that can be accurately modeled as, $y = d\delta + \varepsilon$. This assumption implies that a fitting outcome should be continuous, as opposed to binary or categorical.

ii. The $K \times K$ matrix $X'X$ is well behaved such that it approaches a finite, existent, non-singular matrix as N increases.

iii $E[d_i'\varepsilon_i] = \mathbf{0}$, such that x_i and u_i are uncorrelated.

iv. The data are independently and identically distributed over $i = 1,2, \ldots N$ with $E[\varepsilon|d] = 0$, and $Var(\varepsilon_i|d) = \sigma_i^2$

These assumptions work together to guarantee consistency of the OLS estimator.

OLS is limited for evaluating life course causal effects. First, there a number of ways in which life course processes may violate the assumptions set forth above. While OLS assumes a constant linear relationship between the independent variables and outcome, many life processes may not behave this way. Hungerford and Solon (1987) re-examine the relationship between education and wages to find evidence that log wages are not a smooth function of education. They report evidence of a "sheepskin" effect where wages increase with the completion of certain educational thresholds. A model that instead assumed a linear relationship between years of completed education and wages would be mis-specified if the data generating process were more reflective of the 'sheepskin' process.

Second, of particular concern for this chapter, we have limited ability to differentiate between causal and non-causal relationships. Possible correlation between observables and unobservables threatens to bias OLS estimates. In life course research there are often a range of potentially important factors that may be omitted from a model. A number of factors may enter a model through the error component, such as political or historical contexts, unobserved preferences and/or constraints, and other omitted variables. If these factors influence the outcome of interest, they may also influence values of X, leading to a correlation that violates the key assumption of selection on observables ($A3$) and may produce a biased estimate of β. Nothing in the standard OLS framework implies that the estimated relationship between the independent and dependent variable are causal. OLS offers a way of estimating the conditional change in one measure associated with a change in a related measure. Any attempt to interpret this change as causal is based on a theory of the mechanisms at play, which extends beyond the mathematical properties of OLS.

OLS tends to be a reasonable starting point for studying continuous measures, but inappropriate for binary and categorical outcomes due to violation of the linearity assumption ($A1$). In such cases, we turn to binary and discrete choice models for inference. Binary choice models are generally formulated from the conceptual framework of studying a latent variable y^*, using observable

$$y = \begin{cases} 1 & y^* \geq 0 \\ 0 & else \end{cases}$$

where $y^* = x_i'\beta + \epsilon$. Given this identity, we can write:

$$\begin{aligned}\Pr[y_i = 1|x_i] &= \Pr(y^* \geq 0) \\ &= \Pr\,(x_i'\beta + \epsilon \geq 0) \\ &= \Pr(x_i'\beta \geq -\epsilon) \\ &= \Pr(x_i'\beta \leq \ \epsilon) \\ &= F(x_i'\beta)\end{aligned}$$

Where $F(\cdot)$ is the cumulative distribution function (CDF) of ϵ. We can transform this into a probit or logit estimator by making appropriate assumptions on the distribution of ϵ, and thus the functional form of $F(\cdot)$. If$\epsilon \sim N(\mu, \sigma^2)$, $F(X\beta) = \Phi\left(\frac{X\beta - u}{\sigma}\right)$, where $\Phi(\cdot)$ denotes the standard normal CDF. Alternatively, if ϵ obeys the CDF of the logistic distribution, $F(X\beta) = \Lambda(X\beta)$ where $\Lambda(\cdot)$ denotes the CDF of a logistic distribution leading to the estimation of a logit model. Multinomial logit models are also available to estimate effects on categorical outcomes with more than two values. ϵ could also follow an extreme value distribution implying complimentary log-log estimation. Linear probability models (OLS on a binary outcome) are also sometimes used, however, the resulting parameter estimates have an expected bias due to the clear violation of the functional form assumptions in $A1$ above. Binary outcome modes are essentially a transformation an OLS model, and for that reason they inherit many of the same limitations faced by OLS models for causal inference over the life course. We face potential uncertainty about the true functional form governing our process of interest, and the uncertainty may increase with binary models based on the assumed structure of a latent variable y^* and its

assumed relationship to observed y_i. Most importantly, we face the same uncertainty concerning the assumed independence of observed x_i and unobserved ϵ_i.

Lagged dependent variable models are a variation of OLS or logistic regression models that utilize data over the life course to obtain better estimates of causal effects. Such models include a pre-treatment measure of the outcome as a regressor, and in so doing help control for differences in pre-treatment characteristics that may bias effect estimates. However, we note a few limitations of this approach. First, the approach requires the availability of a pre-treatment outcome measure, which may not be available in some research settings. Second, controlling for pre-treatment differences does not control for other changes that occur between pre-treatment measure and follow-up. This implies a potential need to control for changes that occur in the interim, and to limit the interim time period to the extent possible to limit the likelihood of unobservable changes influencing effect estimates. Lagged dependent variable models and fixed effects model are similar in that they both incorporate repeated measures of an outcome. They differ in that while fixed effects models assign special status to pre-treatment outcomes, by incorporating them into the dependent variable, lagged dependent variable models treat outcome observations from previous period(s) as simply another regressor in explaining the level of the post-treatment outcome.

B. Matching Models

Matching estimators are used to estimate treatment effects by taking the average difference in a selected outcome between individuals with the same pre-treatment observables and different treatment assignments. The fundamental difficulty in estimating treatment effects

involves the impossibility of observing the same observation under observed and counterfactual conditions. Given an outcome y that depends upon observables X, unobservables ε, and treatment status d, we can estimate an average treatment effect by

$$\delta = E[y|X_i, \varepsilon_i, d = 1] - E\left[y\middle|X_j, \varepsilon_j, d = 0\right]$$

where $\mathrm{E}(\cdot)$ is the expectation operator, $X_\mathrm{i} = X_\mathrm{j}$, and $\varepsilon_\mathrm{i} = \varepsilon_\mathrm{j}$. The above estimates δ by computing the expected difference in outcome value among people who differ only by treatment status. Given that individuals have identical profiles for observables X, we attribute average differences in values to differences in treatment assignment. This assumes that conditioning on X is sufficient for achieving conditional independence of the outcome and the treatment such that $Pr(d = 1|X) \perp\!\!\!\perp \mathrm{y}$. Matching methods provide a conceptual framework that parallels experimental settings for evaluating estimated effects in observational settings where assignment to treatment is non-random.

Matching estimators may introduce a dimensionality problem that strains available data. Given the need to condition up $X \in R^k$, where R^k R denotes the real numbers and k indexes dimensionality. the data needs to have a sufficient number of $K-$dimensional matches to facilitate estimation of an expected difference. If K is relatively large the data may not have the available matches. Even for small values of K where each dimension is continuous, having available matches may be problematic. Rosenbaum and Rubin (1983) recommend propensity score matching to reduce the dimensionality of the matching problem. Given that assignment to treatment depends upon observables W, we model the probability treatment as

$$Pr(d = 1|X) = F(X\Gamma)$$

where $F(\cdot)$ is a cumulative distribution function, and a binary model is estimated, usually probit or logit. Estimating such a model allows the calculation of $Pr(\widehat{T = 1}|X)$, a one dimensional measure that summarizes one's likelihood of treatment given available observables. We then match treated and untreated observations based on the one-dimensional estimated probability of treatment, instead of matching on K-dimensional observables. This approach facilitates average effect estimates with smaller datasets, while permitting more measures to be considered in the matching process.

The above discussion of propensity scores and matching is highly non-technical and omits the mathematical assumptions that validate the method. See Rosenbaum and Rubin (1983, 1985) for a more formal discussion of the method and see Caliendo, Marco, and Kopeinig (2008), Leuven and Sianesi (2014), and Morgan and Harding (2006) for practical guidance on implementing the method. There are many examples of propensity score matching used as a tool to estimate treatment effects in social research over the life course. See Brand and Halaby (2006) and Brand, Pfeffer, and Goldrick-Rab (2014) for a few recent examples of research on the effects of higher education on life outcomes.

Matching estimators offer an approach for effect estimation when parametric regressions assumptions may be violated and experimental approaches may be unavailable. However, like regression approaches, matching models to estimate causal effects over the life course are limited. When using a propensity score matching approach (or a regression approach), one needs a fairly strong model of the treatment assignment process of interest. With some natural processes, it may be extremely difficult to construct a model of endogenous treatment

assignment. We must again make assumptions concerning unobservables, ε. Though we often expect individuals with similar observables to have similar unobservables, this may not be the case. Given a life course perspective where ε encompasses all prior characteristics about an observation not captured by X, significant differences between ε_i and ε_j, for $i \neq j$ are conceivable and at times expected.

C. Instrumental Variable Models

In many research situations, OLS effect estimates are biased due to a violation of the ignorability, or selection on observables, assumption iii above. The following heuristic treatment demonstrates the bias in OLS relative to an IV approach. Suppose we wish to estimate the effect of d_i on y_i, given knowledge that an unobservable, z_i, exists that is correlated with d_i and affects y_i. OLS has no way of distinguishing between the effects on d_i on y_i, and the correlated effects of z_i on y_i. IV estimation addresses this limitation of OLS by introducing an instrumental variable, w_i, that is correlated with x_i, but uncorrelated with both w_i and u_i. Mathematically, this implies that assumptions 1 through 3 above hold. Intuitively, it implies the existence of a measure that shifts d_i independently of both z_i and ε_i, to produce an estimate of δ that is unbiased. The bivariate IV estimator takes the form:

$$\hat{\delta}_{IV} = \frac{\frac{1}{N}\sum_{i=1}^{N} w_i y_i}{\frac{1}{N}\sum_{i=1}^{N} w_i d_i}$$

where w_i is an instrument. The instrumental variable, w_i, satisfy the following assumptions:

$i.$ w_i must be correlated with d_i such that $E[x_i z_i] \neq 0$.

ii. w must be uncorrelated with z_i such that $E[w_i z_i] = 0$.

$iii.$ w_i must be uncorrelated with ε_i such that $E[w_i \varepsilon_i] = 0$.

See Cameron and Trivedi (2005, 2010, chap. 6), Greene (2012, chapter 8) and Wooldridge (2010, 2013) for a more general treatment of IV estimators.

In life course research, it is extremely difficult to identify an IV that satisfies the assumptions above. It is typically easy to find a measure that is correlated with the treatment of interest, x_i, but difficult to argue that a potential instrument is uncorrelated with an unobserved confounder, and that the potential instrument has no independent effect on the outcome of interest, y_i, beyond its effect through x_i. Lleras-Muney (2005) offers an example of identification of an IV in life course research. She aims to estimate the causal effect of education on adult mortality using a method that accounts for the likely existence of unobservables that may lead individuals who choose to obtain greater education to make other choices that affect their mortality, a phenomenon that would lead to a bias in the estimated affect of education on mortality. Lleras-Muney (2005) employs an IV approach using changes in compulsory state mandated schooling levels as an instrumental variable. Changes in compulsory education laws increase the education level of students who would have discontinued their education earlier in the absence of the policy constraint.

D. Fixed Effects Models

Individual fixed effects models are useful when data contain repeated observations from a given unit, more likely when we estimate effects over the life course, and each unit is expected to

have a time invariant value, α_i, that shifts the outcome of interest, y_i, in addition to changes in y_i that are correlated with X. This yields a model:

$$y_{it} = \alpha_i + X_{it}\beta + \varepsilon_{\mathrm{it}}$$

where i denotes a unit of observation and t denotes the time period of the observation. In this model, OLS leads to biased estimates of our parameter of interest, β, if α_i is correlated with X_{it} such that $E[d_{it}'\alpha_i] \neq 0$. A fixed effects model produces an estimate of β unbiased by α_i. Fixed effect's models avoid a bias from α_i by estimating $\hat{\beta}$ using within person/unit variation in X_{it} and y_{it}. We employ this approach by demeaning values of $y_{it}, \alpha_i, X_{it}, \varepsilon_{it}$, and estimating $\widehat{\beta_{FE}}$ using an OLS regression of the deviations in X_{it} on the deviations in y_{it}. Since $\overline{y_{\imath t}} = \overline{\alpha_\imath} + \overline{d_{\imath t}}\beta + \overline{\varepsilon_{\mathrm{it}}}$, we can subtract $\overline{y_{\imath t}}$ from both sides of D.1 and rearrange terms to get:

$$y_{it} - \overline{y_{\imath t}} = \alpha_i + X_{it}\beta + \varepsilon_{\mathrm{it}} - \overline{y_{\imath t}}$$

$$y_{it} - \overline{y_{\imath t}} = \alpha_i - \overline{\alpha_\imath} + (X_{it} - \overline{X_{\imath t}})\beta + \varepsilon_{\mathrm{it}} - \overline{\varepsilon_{\mathrm{it}}}$$

$$\tilde{y}_{it} = \tilde{X}_{it}\beta + \tilde{\varepsilon}_{it}$$

Where $\tilde{y}_{it} = y_{it} - \overline{y_{\imath t}}$, $\tilde{X}_{it} = X_{it} - \overline{X_{\imath t}}$, and $\tilde{\varepsilon}_{it} = \varepsilon_{\mathrm{it}} - \overline{\varepsilon_{\mathrm{it}}}$. Note that $\alpha_i - \overline{\alpha_\imath} = 0$ since α_i is time-invariant, which leaves our fixed effects estimate, $\hat{\beta}_{FE,}$ α_i free. $\hat{\beta}_{FE}$ is an unbiased estimator of β if $\tilde{X}_{it}$ and $\tilde{\epsilon}_{it}$ are uncorrelated. That is, time-varying processes that may threaten the plausibility of the ignorability assumption remain a concern. Budig and England (2001) offer an example of individual fixed effects models, applied to an analysis of the wage penalty of motherhood. They found a wage penalty of approximately 5 percent per child, controlling for marital status and human capital variables.

Sibling models are a subset of fixed effects models where fixed effects are assumed to be family specific instead of individual specific. Such models are widely used to understand the effects of life events on children in an analytical framework that controls for differences between children and families that may otherwise bias effect estimates. Currie and Thomas offer an example of sibling fixed effects models, applied to an analysis of the long-term effects of Head Start on students' academic achievement. They find significant gains in test score and declines in the likelihood students will repeat a grade, but the positive benefits are concentrated among white children.

One potential concern with fixed effects model concerns the ambiguity of exactly what is subsumed in the fixed effects parameter. It is not always clear that a fixed effects model is not introducing an endogenous selection bias, and possibly dampening the treatment effects of interest. Researchers should carefully consider whether the differencing process undermines the variability they wish to explain.

E. Experimental and Quasi-Experimental designs

Experimental research designs help facilitate identification of treatment effects by maintaining full control over the treatment assignment process. Knowledge of the treatment assignment process allows a stronger argument for claiming that treatment is uncorrelated with unobservables that may lead to bias in parameter estimates and that there are no selection effects governing who receives treatment, which implies no systematic differences between the treated and untreated.

Randomized Control Trials (RCTs) are the most common type of experiment. In RCTs each unit of study is randomly assigned to either a treatment or control group. In many RCTs, the researcher obtains a baseline measure of the outcome of interest at or near the time of randomization. Then, one group receives the treatment, and then a follow-up measure is taken that captures treatment effects. Effect estimates may be calculated via a difference calculation that looks at the incremental change in the outcome measure of the treatment group relative to the control group, or regression analysis which can further control for differences in observables among the treatment and control groups that may randomly occur. While RCTs are highly effective at reaching causal effect estimates, they are often very expensive. For many questions concerning life processes, an experiment mandating that individuals endure certain treatments is precluded for both practical and ethical reasons. Though RCTs are designed to offer clean effects, they do not always address the intuitive question of interest. RCTs indicate the effect of being assigned to a particular treatment group. In many interventions, there are additional considerations of a take-up rate and dosage. At times, experimental results indicate a zero effect potentially because the parameters of program design allow treated individuals to receive less than the recommended treatment dosage. Those engaging in experimental design try to safeguard against opportunities for participants to make endogenous choices that affect program receipt and effect estimates.

Quasi-experiments aim to achieve the benefits of a full experiment in a setting where the researcher does not have full control over the treatment assignment process. One of the most well known quasi-experiments in the stratification literature is the Moving to Opportunity (MTO)

experiment. This intervention randomly assigned two types of housing vouchers to low income families and in order to assess neighborhood effects on economic, health, and child outcomes. One voucher afforded families the option of moving to another low-income neighborhood, and the other offered the option of moving without constraint. Each treatment family chose whether or not to exercise their option to move, and this lead to selection issues and controversy concerning the validity evaluation results [see Goerhrig (2003), Ludwig et al. (2008), and Sanbonmatsu, Lisa, et al (2011) for further discussion of MTO].

A natural experiment is a type of quasi-experiment where the researcher has no control over the treatment assignment process, but has reason to believe that the natural treatment process yields (1) and (2) above. Sharkey (2010) examines the acute effects of local homicides on cognitive performance of children. He uses the exogeneity between the timing of local homicides and the timing of vocabulary and reading assessments in Chicago schools to craft a natural experiment that gives estimates of the short-term decrease in test scores associated with exposure to local homicide. He finds significant negative effects on tests administered up to seven days after a local homicide. The acute affect weakens over time and becomes arbitrarily close to zero on after four weeks.

3 Causal Inference in the Presence of Heterogeneity, Time-Variation, and Mediation

A. Identifying Heterogeneous Causal Effects

If there is treatment effect heterogeneity, average treatment effects can vary widely depending on the population composition of the treated and thus simple averages do not have a straightforward interpretation. Indeed, an important development of the causal inference literature is the recognition that treatment effects are likely to be heterogeneous (Angrist and Krueger 1999; Brand and Simon-Thomas 2013; Heckman, Urzua, and Vytlacil 2006; Morgan and Winship 2012, Xie, Brand, and Jann 2011; Xie 2011). This kind of heterogeneity does not merely reflect group differences at the baseline that can be "controlled for" by covariates in regression or matching models, or fixed effects. The recognition that treatment effects may vary by the probability of treatment, beyond response variation by selected covariates like gender or race, has led to new methods of causal inference and to refined interpretations of effect estimates derived from existing methods (Brand and Xie 2010; Elwert and Winship 2011; Morgan and Todd 2008; Morgan and Winship 2012; Xie 2011; Xie, Brand, and Jann 2013). Despite widespread belief by practitioners, traditional regression estimates do not represent straightforward averages of individual-level causal effects if individual-level variation in the causal effect of interest is not random. Instead, they give a peculiar type of average – a conditional variance weighted average of the heterogeneous individual-level effects, where population composition weights can produce widely different effect estimates.

Regression and matching models can be used to recover subpopulation treatment effects of interest, including the treatment effect on the treated (*TT*) and the treated effect on the untreated (*TUT*). Let us define the average difference among those individuals who were actually treated, the *TT*:

$$\overline{\delta}_{TT} = E(y^1 - y^0 \mid d = 1),$$

and the average difference among those individuals who were not treated, the *TUT*:

$$\overline{\delta}_{TUT} = E(y^1 - y^0 \mid d = 0).$$

Statistical modeling to explore empirical patterns of effect heterogeneity as a function of the propensity score have also been employed to recover patterns of treatment effect heterogeneity (Xie, Brand, and Jann 2013). Several recent studies have adopted this approach to address questions involving heterogeneous effects of higher education on a range of life course outcomes (Brand 2010; Brand and Davis 2011; Brand, Pfeffer, and Goldrick-Rab 2014; Brand and Xie 2010; Musick, Brand, and Davis 2012). Instrumental variable models, in the presence of effect heterogeneity, may be interpreted as identifying local average treatment effects (*LATE*), those effects corresponding to subpopulation on the margin of treatment participation, induced by the particular instrument under consideration.

B. Identifying Causal Effects with Time-Varying Treatments and Time-Varying Outcomes

Life course research often involves the analysis of effects of events that occur over time, which raises particular concerns with causal inference. Individuals who experience the event of interest early on may do so for different reasons than those who experience the event later. Researchers must carefully attend to the conceptual and theoretical issues underlying life course treatments. For example, Brand and Simon-Thomas (2014) look at the effects of maternal job displacement on educational and social-psychological outcomes of children using a propensity score model estimation framework. Correcting for selection into displacement requires a

reasonably strong model of which children are most likely to experience maternal job loss. Models that predict the likelihood of experiencing displacement using only covariates available at or before the child's birth do a much better job of predicting displacement events that occur in early childhood relative to displacements experienced in middle childhood and adolescence. This illustrates some of the difficulty one may encounter when using one model to explain an occurrence that happens at different times for different reasons. However, Brand and Simon-Thomas (2014) also partition maternal displacement into three periods across childhood, and assess effects that occur at varying points in young adulthood, adopting the conceptual framework of Brand and Xie (2007) we discuss below.

Brand and Xie (2007) discuss some of the conceptual challenges involved with estimating causal effects of non-repeatable and non-reversible treatments that occur at different points in time, and that affect outcomes that are measured at different points in time. Table 1 shows the possible combinations of time varying treatments and outcomes. We also list the effect of interest in each case for some outcome y_i measured at time v. Treatment of study occurs at time d, and we operate in an environment where non-repeatable and non-reversible time varying treatments may occur at multiple times before and after the treatment event of interest. In the table below, we care to study the treatment events that occur at time t, but there may also be observations who experience the event at time t^+. Table 1 shows four possible ways in which time-varying treatments and outcomes may interact.

Case 1 considers the effects of a time-invariant treatment on a time-invariant outcome. This is the classic assumed two period framework where treatment occurs in some period and an

outcome is realized thereafter. Case 2 consider the effect of a time-invariant treatment on a time-varying outcome. This would be useful if we care to assess how a disruption affects an outcome measured in repeated future periods. These types of analyses help determine how treatment and time since treatment interact in the determination of causal effects, and offer evidence as to whether affects grow or weaken over time. Case 3 considers a time-varying treatment and time-invariant outcome. This model can shed further light on how processes unfold over time. For example, one could use this approach to understand whether experiencing marital disruption has dramatically different implications for a women's earnings at age 40 depending upon whether she's in her 20's or 30's when the disruption occurs. Case 4 considers time-varying treatment and time-varying outcomes. For example this set-up could assess the differences in the effects of job loss at age 20 and age 25 on earnings at age 30 and age 35 relative to the effect of job loss at age 30 and age 35 on earnings at age 40 and age 45. In case 1 the relevant counterfactual to receiving a treatment in period 1 is clearly the subsequent outcome that would occur in the absence of period 1 treatment. The same holds for case 2, except that we can employ measures from different post-treatment periods as our outcomes of interest. The counterfactual becomes slightly more complex in cases 3 and 4. Whereas cases 1 and 2 segment the sample based on single period receipt, cases 3 and 4 must consider compare a treated group to a counterfactual of not receiving treatment within the data's observation period and all future periods through outcome measurement.

Sample and cell size consideration affect the feasibility of producing estimates. Comparisons are made between a subgroup who experience treatment at time t, and the

subsample who has not experienced treatment up to time t. This implies that all individuals who experienced treatment before time t are excluded from the stated effect estimate. When we care about assessing the difference associated with experiencing treatment in different periods, we need enough individuals experiencing treatment within those periods to produce reliable estimates. A wider interval benefits from the inclusion of more observations, which can yield more precise estimates, but means that our pre-treatment covariates have potentially reduced ability to predict the probability of treatment exposure.

B. Identifying Causal Effects with Mediating Mechanisms

In social research aimed at causal inference over the life course, researchers often aim to assess which mediating mechanisms transmit the effect from treatment to outcome. Too often researchers simply add an additional intermediary variable to the model, and then assess the degree to which the effect of the treatment on the outcome has changed in response to the inclusion of the additional variable. This type of analysis, even when a great deal of attention ahs been paid to accounting for selection into treatment, often fails to attend to the causal process relating the treatment to the mediating variable or the mediating variable to the outcome. We argue that if researchers aim to asses indirect causal effects, they should devote the same attention to causal processes linking the mediating mechanisms as they do to the primary treatment of interest.

Judea Pearl (2009) and colleagues have developed directed acyclic graphs (DAGs), graphical tools that are becoming increasingly popular by providing a visually tractable framework for assessing whether a model is identified and the mechanisms that may mediate effects of interest (Elwert 2013). Conventions include that an arrow indicates a direct causal effect ($d \rightarrow y$), while a missing arrow indicates no causal effects ($d\ y$). A line indicates two variables are endogenously correlated with no causal direction ($x_1 - x_2$). A causal path can be depicted by $d \rightarrow y_1 \rightarrow y_2$, $x \rightarrow y_1 \leftarrow d \rightarrow y_2$ is a non-causal path from x to y_2. A variable with two arrows along the path pointing into is a collider. For example, y_1 is a collider along $x \rightarrow y_1 \leftarrow d \rightarrow y_2$. We open our estimation to endogenous selection bias when conditioning on a collider variable. Encoded within the DAGs are rules for moving from causation to association. Chains represent causal associations ($d \rightarrow y_1 \rightarrow y_2$), forks represent confounding ($d \leftarrow z \rightarrow y_1$) and inverted forks represent endogenous selection ($d \rightarrow y_1 \leftarrow y_2$). We represent the causal process with a DAG in our empirical example below.

3 Research Question, Data Generating Process, and Causal Estimation

A. Description of the Data and the Data Generating Process

For this chapter, we have created simulated longitudinal data that follows adolescents through four years of secondary education. The data were collected for two cohorts of young people who enter high school two decades apart. Cohort one begins in 1980 and cohort 2 begins in the year 2000. The data for each cohort include roughly 5,250 families who have over 7,000 high school aged children. Each family has between one and three children in the data. The data

contain standard demographic characteristics for all children and families. Race and ethnicity classifications include white, black, Asian, and Hispanic. Family structure classifications include a child/children living with both parents married, unmarried single mother, or living with an unmarried single mother with a father/father figure. Single father headed households and complex non-traditional household structures were excluded. Parent's education indicates whether each parent's highest achieved education is less than a high school degree, a high school degree, a college degree, or a graduate/professional degree. Families live in one of five neighborhoods that differ in their demographic characteristics and risk rate of exposure to the treatment of interest. For example, respondents in some neighborhoods have a 3 percent chance of random exposure to violence, whereas the probability of random exposure in other neighborhoods is as high as 12 percent.

The simulated data were generated with the goal of replicating the complexities of survey data in a framework that allows for accurate estimation of counterfactual outcomes. To this end, some variables in the data were constructed to produce joint distributions similar to those that we tend to observe in U.S. data, while other measures were specified subjectively. The data generating process begins by specifying the demographic traits of a family. Families are characterized by race, household income, community of residents, mother's education, father's education, family structure, and the number of high school aged children who serve as respondents in the data. For each family, one of four races was randomly assigned to yield a sample that is approximately 11% Asian, 25% black, 25% Hispanic, and 40% white. Conditional upon race, a household income value is chosen from race specific distributions. The income

distribution for the year 2000 cohort was defined to approximate the race specific U.S. income distribution reflected in March 2009 Current Population Survey data. The 1980s income distribution is a transformation of the 2000s distribution where the 2000s income distribution first order stochastically dominates the 1980s distribution.[1] Household structure is randomly assigned according to an income quintile specific distribution. Available structures include (1) household lead by married parents, (2) single mother headed household, (3) household lead by cohabiting parents. In both cohorts, households with higher income are more likely to be lead by married parents, while households with lower incomes are more likely to have to have a single head, or cohabiting parent(s). Families from the 1980s cohort are more likely to be married while families in the 2000s have more single mother lead households and households with cohabiting parents. Conditional upon race and household structure, each family is randomly assigned between one and three high school aged children. The same transition probabilities for the fertility process apply to both cohorts.

Parental education takes one of four values (less than high school, high school degree, college degree, and graduate/professional degree), and was assigned based on income draws. Families with higher household income receive parental education draws from distributions that place higher probabilities on higher education levels. Education levels of both mother and father are drawn from the same distributions. This approach creates an expected correlation between parent's education and household income, as well as educational homogamy between parents.

[1] The 1980s income distribution has a smaller support (with a maximum value of $150,000 compared to $250,000 for the 2000 distribution) and a lower mean value. We implement these distributions by drawing a random integer from a uniform distribution on the closed interval [0, 10000], dividing the integer by 100 to get a percentile, and mapping that percentile to an income value. See Appendix Table A for the mapping of percentiles to income values.

Last, five communities are determined that differ in terms of racial composition and probability of exposure to violence. Community 0 has a neutral racial composition that mirrors the sample wide racial makeup. All other communities have a proportionally dominant racial group.[2] The demographic characteristics of families were designed to capture many of the issues we observe in survey data. We generally see correlations between education, income, race, and other measures that obfuscate the direct relationships between demographic measures and an outcome of interest. In collected data with unobserved counterfactuals, we have no way of knowing the extent to which these correlations bias estimates of interest. In these simulated data, we can observe counterfactual outcomes, and thus calculate the difference between regression estimates and actual counterfactual outcomes.

The characteristics described in Table 2 in the two sample periods (1980-84 and 2000-04) have notable differences. First, children are much more likely to live with two married parents in the 1980s, while mother-headed households and households with cohabiting parents are relatively more common in the 2000s. Mothers, however, are significantly more likely to have a graduate or professional degrees in the 2000s relative to the 1980s. All but one community became less segregated over the 20-year time period with the largest racial group representing a smaller percentage of the population of each community. We have academic data for students for four years of high school. Measures include whether the student enrolled in high school during each year, how many credits were earned towards graduation, and whether the student earned enough credits to graduate after four years. We do not follow students' trajectories beyond this

[2] See Appendix A for further details on the constructed data.

time frame. The data also includes labor market participation for students, including whether students worked in the labor market in a given year, what proportion of their time was allocated to wage labor, and their wage rate.

B. Treatment Specification and Decision Mechanisms

The treatment of interest is exposure to neighborhood gun violence (ETV), and the outcomes include the likelihood of high school completion and the number of credits earned towards a high school diploma in a 4-year period. Receiving the "treatment" in these data is equivalent to answering the following question affirmatively: In the previous school year, have you seen anyone shot or shat at, or has anyone close to you been shot or shat at? Among those who answer this question affirmatively, it is unknown whether the respondent was exposed to fatal or non-fatal shooting. It is also unknown whether the respondent was exposed to more than one qualifying event. It is, however, known that fatal shootings are rare and exposure to multiple shootings in a given year is uncommon. Patterns of gun violence across neighborhoods have changed very little over the 20-year time period. Incidents are relatively uncommon in neighborhoods one and two, while neighborhoods three through five tend to experience violent occurrences more frequently.

Community-level graduation rates are negatively correlated with incidents of gun violence. While this is consistent with an inverse relationship between exposure to violence (ETV) and academic success, there are a range of other differences between communities that could partially or fully explain the differences in graduation rates. These include observable

differences in parental education, household income, school quality, and racial composition, and unobservable differences in ability, personal valuations on education, and varying opportunity costs. We aim to assess whether exposure to violence has a causal effect on high school completion and credits earned towards graduation. The following section uses some common causal inference methods used in life course research to address these questions.

The ETV process has modest differences between the earlier and later cohorts. In the 1980s data, exposure rates are neighborhood dependent but otherwise random. Residents of each neighborhood have a fixed year-specific probability of community violence exposure, and individual characteristics play no part in determining who experiences the event of interest. The violence exposure in the 2000 cohort is similar, except individuals with a high preference for leisure are more likely to experience exposure. This adds endogeneity to the process governing the receipt of treatment, and this is an issue we often observe in survey data.

Educational outcomes are the result of a sequence of static decision problems. Simulated respondents make decisions in 4 separate periods, representing the four years of a traditional high school career. Given individual preferences for consumption and leisure, as well as a wage offer, an endowment/financial allowance from the household, and individual valuation of high school graduation, simulated respondents choose how to allocate a fixed allotment of time between academic labor, wage labor, and leisure time. Income from each period is used to finance consumption for that period only. Inter-temporal savings are not permitted, and simulated respondents solve a sequence of intra-temporal problems instead of solving a single inter-temporal problem. It might be more realistic to model this decision inter-temporally,

however, it was our objective to write the simplest possible model that would yield non-trivial results. This approach succeeds at doing that.

In this framework, high school is modeled as a four period process where students must earn 12 academic credits to graduate and respondents always aim to divide credits evenly over their remaining academic years. Academic credit in period t, c_t, is the product of academic labor, l_t^a and ability, a_t such that

$$c_t = l_t^a \cdot a_t$$

Respondents may have low, medium, or high ability, corresponding to $a_t \in \{9, 14, 20\}$, respectively.[3] Solving for l_t^a in (1) gives

$$l_t^a = \frac{c_t}{a_t}$$

showing that kids with higher ability require less labor to earn a fixed amount of academic credit. Respondents use the labor demand estimate in (2) to facilitate a decision problem concerning whether to continue their education, in period t.

Simulated respondents have preferences over consumption, leisure, and high school completion. In each period they make time allocation decisions that determine their consumption, leisure, and academic credit earned for the period. Time can be allocated to any combination of academic labor, l_t^a, wage labor, l_t^w, and leisure, $1 - l_t^a - l_t^w$. Simulated respondents determine how they would divide their time if they enroll (do not enroll) in school,

[3] These values were chosen to simulate ability heterogeneity, and generate differences in the amount of academic labor required to complete a year of schooling. The values have no deeper meaning beyond serving that purpose.

and choose the option with the highest expected payoff. See Appendix A for details on the simulation process.

C. Treatment Effects

Exposure to violence (ETV) registers as a traumatic experience that depresses productivity and erodes gains from leisure time. This is modeled as a loss of η units of time that would have otherwise gone towards labor or leisure. This leaves the agent with $1-\eta$ units of time to allocate over leisure, academic labor, and wage labor. ETV also lowers the perceived gains to a high school degree by raising the subjective probability that one may not survive to reap the returns to academic labor. This is modeled as a decrease in the graduate return, G, of ν units. Both mechanisms leading to ETV effects can lower graduation rates. The η penalty can lower available time to the point that supplying the often desired $\frac{c_{t,i}}{a_i}$ units of academic labor is no longer feasible or no longer optimal. The decrease in G can also make enrollment suboptimal if it moves the value of continuing below a certain threshold. Both penalties should lead to heterogeneous non-linear effects on high school completion and credits earned since the simulated agent's actions post ETV are also highly affected by other agent-specific parameters.

5 Effects Estimates

A. True Causal Effects

In the 1980s (2000s) data, 15.8 percent (17.5 percent) of the simulated sample experienced some exposure to violence over the 4-year period. Of those exposed, 93.0 percent (89.8 percent) of the 1980s (2000s) sample experienced ETV event(s) in exactly one of the four year intervals. Of the 1980s (2000s) sample, 6.5 percent (9.6 perent) experienced ETV event(s) in 2 of the 4 years, and the remaining percentage experienced ETV events in 3 of the 4 years. Figure 2 summarizes the distribution of ETV exposure over simulated high school years. In both samples we observe a saw-tooth pattern, where exposure falls in year 2, rises dramatically in year 3, and then falls slightly in year 4. Overall, exposure rates are uniformly higher in the 2000s data.

The simulated environment allows the estimation of true effects by comparing observed outcomes with observable counterfactual outcomes. For our sample of simulated respondents who were exposed to violence, we calculate the true average effect on an outcome as

$$\Delta y_{ETV} = \bar{y}_{ETV} - \bar{y}_{\sim ETV}.$$

We also calculate the average percent change in the outcome as

$$\%\Delta y_{ETV} = \frac{\bar{y}_{ETV} - \bar{y}_{\sim ETV}}{\bar{y}_{\sim ETV}}.$$

We wish to examine four measures: (1) periods employed during high school; (2) total time devoted to wage labor; (3) total time devoted to academic labor; and (4) high school completion.

Tables 4a and 4b summarize the calculated effects for both cohorts. In the 2000 cohort data, we observe negative ETV effects on all listed outcomes. ETV leads to a decrease in years

of work and hours worked, as well as a decrease in time devoted to academic labor and a lower proportion of students who graduated in four years. We see generally similar findings for the 1980s data with some differences. First, the full 1980s cohort works all four years of high school regardless of ETV. This yields no effect on employment. We do, however, see a decrease in hours, implying the sample continues to work, but may supply fewer hours of wage labor given ETV exposure. Relative to the 2000 cohort, ETV leads to a much greater decrease in academic labor hours, and a greater decrease in graduation rates.

Simulated agents compare the expected return of enrolling in school versus not enrolling in school each simulated year, and choose the option with the highest value. Simulated agents in borderline regions will not enroll if they experience ETV in a given period, but will enroll otherwise. Roughly 42 percent of the sample falls in the borderline region for the 1980s cohort, while only 23 percent of the sample is borderline for the 2000s cohort. In addition to these distributional differences, the 1980s cohort suffers from a selection problem such that simulated agents who are already at risk of dropping out are roughly 20% more likely to experience ETV. These differences lead to the noticeable differences in effects between cohorts. The analysis shows decreases in all outcome measures, with the exception of periods worked for the 1980s cohort. Below we use methods discussed above to attempt to recover the same information disclosed by the causal analysis presented here.

B. Causal Effect Estimates

This section employs methods described above to infer the causal effect of ETV on a set of outcomes for each cohort. The true counterfactual effects are also listed and compared to the effect estimates from various methods. True counterfactual effects on observables differ between cohorts as a result of differences in the distributions of latent and demographic variables. Tables 5a and 5b show controlled and uncontrolled regression coefficients, whether the true effect falls within the 95 percent confidence interval of the effect estimate, and whether the effect estimate has the same sign as the true effect. Controls include mother's education, father's education, race, respondent wage offer, household structure, parental income, and community of residence. There is no temporal variation in the control variables.

Estimates from the employed methods generally had the same sign as the true effect. The instrumental variables estimates were the only exception. Four of the 8 instrumental variables estimates (where the instrument was an exogenous policy enacted that lead to a sharp increase or decrease in the likelihood of experiencing ETV) were positive when the true effect was negative. There were mixed results concerning the success with which a 95 percent confidence interval held the true effect. OLS and propensity score matching estimates appear to be most successful at generating a 95 percent confidence interval that held true effect estimates. The point estimates from fixed effects estimates often appeared to be fairly close to the true parameter values, but in this case the smaller standard errors from fixed effects estimation lead to very small confidence intervals that often excluded the true values. An informal assessment of estimates from this simulation may lead one to conclude that OLS, Propensity Score Matching, and fixed effects

estimates were most successful at approximating the true causal effect of ETV on the listed outcomes, while instrumental variables estimators were least reliable.

Based on these estimates, we conclude that ETV leads to negative shocks to the examined outcomes. In the year following a ETV, simulated respondents are less likely to enroll in school. Over the course of the high school career, ETV is associated with less overall time devoted to academic and wage labor, and a lower probability of completing high school in four years. There is no effect of ETV on the number of high school years in which a simulated respondent was employed in the 1980s data, but an apparent negative effect is present in the 2000s simulated data. Other estimates also tend to vary between cohorts. While ETV likely causes negative change in the outcomes studied, we have limited accuracy with which we can estimate the magnitude of the associated changes.[4]

In applied life course research, we generally do not fully understand the role of differences in the distribution of observable and unobservable variables, as well as differences in the treatment selection process over time. In this simulation, the effect of ETV on college graduation was 13 times larger in the 2000s simulated cohort than in the 1980s simulated cohort. There are noticeable differences in the magnitude of true effects on all outcomes between cohorts. It is possible that the distribution of household income and respondent wage offers (which differ across the two cohorts) are largely responsible for differences in the effect of ETV on periods employed. Other differences between the cohorts may be responsible for the

[4] For simplicity, we do not consider heterogeneity in effects here. We assume that effects are homogenous. In applied life course, research, however, researchers should routinely question the underlying homogeneity assumption.

difference in estimates. In the 1980s data, ETV is a completely random process that is independent of individual characteristics. In the 2000s data, ETV is slightly more likely to happen to simulated respondents who have a stronger taste for leisure than consumption (i.e. $\beta_i > \alpha_i$). There was also an overall shift in unobservable preference parameters between the periods. The 1980s cohort has a stronger preference for leisure, while the 2000s cohort has a higher preference for consumption. In this model a taste for consumption is more conducive to continued enrollment and eventual graduation. The unobservable shift in personal preferences and the selection mechanism in the 2000s data may explain a large part of the difference in sensitivity of the educational tract to ETV exposure between cohorts.[5]

C. Mediating Mechanisms

We often care about both the effect of a treatment and the mechanism mediating that effect. The data were generated such that ETV leads to a lower expected return to education (operationalized as a change in a latent variable denoting the personal valuation placed on graduating), and a decrease in the amount of productive time available for academic and wage labor. A lower valuation on graduation implies a lower likelihood of enrolling in school in any period, and thus a lower likelihood of graduating in four periods/years. The lost time may have the indirect effect of forcing a simulated respondent to choose between employment and

[5] We do not assess how exposure to the treatment over time produces variation in individual effects. It is likely that ETV that occurs early in childhood, relative to middle childhood, relative to adolescence might influence the effects on our outcomes of interest.

enrollment. In this simulated example, effects are mediated by constraints and latent measures that are unobservable to the researcher. Fully recovering these mechanisms is likely impossible with only the observable measures in the data. It is also often the case in life course research that multiple mechanisms may work independently or jointly to produce a net effect of interest. The causal DAG presented in Figure 1 depicts the effect of ETV on the probability of high school completion. We enclosed unobserved variables in ovals. There is a correlation between experiencing ETV and the latent value of high school completion, both of which are correlated with neighborhood and family characteristics. Conditioning on neighborhood and family characteristics implies conditional independence between ETV exposure and latent personal valuations on high school completion. Controlling for neighborhood and family characteristics allows us to view the variation in ETV as exogenous, and identify the causal effect of ETV on enrollment decisions and graduation. This estimate would include effects of all mediating mechanisms. In these simulated data, we speculate that the mediating mechanism might involve the negative effect on both labor variables, such that ETV leads to a short term drop in productivity that results in a decrease in the amount of time devoted to academic and wage labor, and a lower likelihood of enrolling in school in the coming academic year. As a result of the productivity drop and decreased likelihood of enrolling, youth fall behind in school and many are unable to make up the difference in time to graduate in the four-year window. This leads to a lower four-year graduation rate among those who have experienced ETV. If all variables below were observable, we could decompose the mechanisms mediating ETV effects by determining what proportion of the effect works through the devaluation of education, and what portion is

attributable to a loss of productive and leisure time. In this case, leisure time and time availability are unobserved. Quantitative researchers are often limited as to which available theories for mediating mechanisms can be empirically tested.

6 Conclusion

This chapter has discussed some of the challenges and strategies underlying life course research aimed at causal inference. Relationships among variables in empirical data alone cannot establish causality. Causal inference requires some knowledge about the data generating process to support an assumption of exogenous variability in our treatment of interest. This may be controlled randomization of the assignment to treatment, exogeneity of a natural treatment assignment process, or conditional independence of treatment assignments after controlling for some set of observables. Research aimed at causal inference often involves trying to understand complex and dynamic processes that depend upon both observed and unobserved factors. We discuss challenges to empirical estimation that arise due to selection effects. We show and discuss how violations of model assumptions may lead to bias in parameter estimates, and discuss common estimation techniques. These include ordinary least squares regression, discrete choice models, propensity score matching models, fixed effects models, and instrumental variables models. Many other potential estimation procedures were omitted from this discussion for simplicity. Additional challenges may include estimating effects with complex counterfactuals and evaluating heterogeneity in effects. Life course research also often involves

assessing how effects vary over time, as well as assessment of mechanisms that intervene and help explain associations between treatments and outcomes.

We generate simulated data that models a complex process over time, and employ the listed techniques to recover the dynamics and effects generated by our manufactured process. Our simulated process considers the effect of exposure to gun violence on the educational outcomes of simulated respondents. We construct a process where simulated respondents have unobserved preferences over consumption, leisure, and high school completion and varying levels of unobserved academic ability. Simulated respondents are characterized by demographic characteristics such as neighborhood of residence, racial classification, household income, and parental education levels that are correlated with their preferences and the monetary resources that determine consumption levels. In each of four simulated periods, respondents decide whether to work in the labor market full-time, work and go to school, or go to school without working. Offering sufficient levels of academic labor over the four-year period results in graduation. Our treatment of interest, exposure to violence (ETV), flags simulated respondents who have seen a person shot or shot at in the previous year.

The treatment has two simulated effects. First, it decreases subjective assessments of personal longevity, which decreases the horizon over which one would expect to reap the returns to education, therefore lowering the overall expected return to education. This is operationalized as a 3-unit decrease in the simulated agent's personal unobserved valuation of graduation. Second, ETV serves as a traumatic experience that depresses productivity in the short run. This is operationalized as a loss of time that could have otherwise been allocated to leisure, academic

labor, or wage labor. The loss of time and devaluation of education together make the choice to enroll in school less likely, which leads to a lower four-year simulated graduation rate for those who experience ETV. This process was simulated for two cohorts. In one cohort, assignment to treatment has an endogenous aspect where individuals who have a stronger taste for leisure than consumption are a percentage point more likely to experience treatment above the ETV exposure rates defined by their neighborhood of residents. This yields a slight selection effect where individuals who are already more likely to dropout are also more likely to experience an event that further encourages early school exit. The cohorts also differ along the distribution of demographic variables that govern familial characteristics.

We aimed to understand the dynamics behind the effect of ETV on 4-year high school completion using the methods described earlier in the chapter. We found that ETV leads to a decrease in the 4-year high school graduation rates. We also find that ETV is associated with a drop in wage and academic labor supply, which lead to the inference that time may be one of the mediating factors transmitting the effect of ETV on 4-year high school completion. We were unable to uncover the second mechanism, which affects graduation due to a decrease in the personal valuation of high school completion. The respondents' simulated valuation on education is an unobserved variable, and for that reason it could not be included in regression models. Estimates for the effects of ETV differed greatly between cohorts even though the coded effect of ETV was the same in both settings. In both simulations, ETV lead to a 3 point decrease in the valuation of education and a 20 percent decrease in time available to allocate between labor and leisure. These direct effects lead to very different indirect effects of ETV on labor

supplies, per period enrolment, and graduation probabilities. We concluded that differences in the distributions of underlying variables lead ETV to have vastly different effects between cohorts.

Of the methods employed, OLS regression and propensity score matching most consistently produce results that were close to the true effects calculated through counterfactual simulation. Fixed effects estimates were also reasonably accurate. Instrumental variables models had the weakest performance. In our simulation, we found that simple linear methods were sufficient for understanding the basic causal dynamics of interest in a highly non-linear model that includes a range of unobserved variables. This lends confidence in our ability to use these tools to correctly discern the dynamics between complex real-world processes, and the underlying observable mechanisms motivating causal relationships. Still, the performance of these estimators in this setting should not be generalized to other settings.

References

Angrist, Joshua D., and Alan B. Krueger. "Empirical strategies in labor economics." *Handbook of labor economics* 3 (1999): 1277-1366.

Brand, Jennie E. 2010. "Civic Returns to Higher Education: A Note on Heterogeneous Effects." *Social Forces* 89(2):417-433.

Brand, Jennie E. and Dwight Davis. 2011. "The Impact of College Education on Fertility: Evidence for Heterogeneous Effects." *Demography* 48(3):863-887.

Brand, Jennie E., Fabian T. Pfeffer, and Sara Goldrick-Rab. "The Community College Effect Revisited: The Importance of Attending to Heterogeneity and Complex Counterfactuals." *Sociological Science* 1 (2014): 448-465.

Brand, Jennie E., and Juli Simon Thomas. "Causal Effect Heterogeneity." *Handbook of Causal Analysis for Social Research*. Springer Netherlands, 2013. 189-213.

Brand, Jennie E., and Yu Xie. "IDENTIFICATION AND ESTIMATION OF CAUSAL EFFECTS WITH TIME-VARYING TREATMENTS AND TIME-VARYING OUTCOMES*." *Sociological methodology* 37.1 (2007): 393-434.

Brand, Jennie E., and Yu Xie. "Who benefits most from college? Evidence for negative selection in heterogeneous economic returns to higher education." *American Sociological Review* 75.2 (2010): 273-302.

Budig and England. 2001. *American Sociological Review*.

Caliendo, Marco, and Sabine Kopeinig. "Some practical guidance for the implementation of propensity score matching." *Journal of economic surveys* 22.1 (2008): 31-72.

Cameron, A. C., and P. K. Trivedi. 2005. Microeconometrics: Methods and Applications. New York: Cambridge University Press.

Currie, Janet and Thomas, Duncan, “Does Head Start Make A Difference” . *American Economic Review*, June 1995, 85 #3, 341-364

Dehejia, Rajeev H., and Sadek Wahba. "Propensity score-matching methods for nonexperimental causal studies." *Review of Economics and statistics* 84.1 (2002): 151-161.

Elwert, Felix. "Graphical causal models." *Handbook of causal analysis for social research*. Springer Netherlands, 2013. 245-273.

Goering, John, Judith D. Feins, and Todd M. Richardson. "What have we learned about housing mobility and poverty deconcentration." *Choosing a better life? Evaluating the Moving to Opportunity social experiment* (2003): 3-36.

Greene, W. H. 2012. Econometric Analysis . 7th ed. Upper Saddle River, NJ: Prentice Hall

Heckman, James J., Sergio Urzua, and Edward Vytlacil. "Understanding instrumental variables in models with essential heterogeneity." *The Review of Economics and Statistics* 88.3 (2006): 389-432.

Hungerford, Thomas, and Gary Solon. "Sheepskin effects in the returns to education." *The review of economics and statistics* (1987): 175-177.

Leuven, Edwin, and Barbara Sianesi. "PSMATCH2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing." *Statistical Software Components* (2014).

Lleras-Muney, Adriana. "The relationship between education and adult mortality in the United States." *The Review of Economic Studies* 72.1 (2005): 189-221.

Ludwig, Jens, et al. "What can we learn about neighborhood effects from the Moving to Opportunity experiment." *American Journal of Sociology* 114.1 (2008): 144-188.

McLanahan, Sara, Laura Tach, and Daniel Schneider. "The causal effects of father absence." *Annual review of sociology* 39 (2013): 399-427.

Mortimer, Jeylan T., and Michael J. Shanahan, eds. Handbook of the life course. Springer, 2006.

Moore, Quinn, Robert G. Wood, Andrew Clarkwest, Alexandra Killewald, and Shannon Monahan. "The Long-Term Effects of Building Strong Families: A Relationship Skills Education Program for Unmarried Parents. Technical Supplement." Princeton, NJ: Mathematica Policy Research, 2012.

Morgan, Stephen and David Harding. 2006. "Matching Estimators of Causal Effects: Prospects and Pitfalls in Theory and Practice." *Sociological Methods and Research* 35(1):3-60.

Morgan, Stephen, and Christopher Winship. "Brining context and variability back into causal analysis." *Oxford handbook of the philosophy of the social sciences* (2012): 319-354.

Morgan, Stephen L., and Jennifer J. Todd. "A diagnostic routine for the detection of consequential heterogeneity of causal effects." *Sociological Methodology* 38.1 (2008): 231-281.

Musick, Kelly, Jennie E. Brand, and Dwight Davis. "Variation in the relationship between education and marriage: Marriage market mismatch?." *Journal of Marriage and Family* 74.1 (2012): 53-69.

Pearl, Judea. "Causality: Models Reasoning, and Inference. Cambridge University Press. (2009).

Rosenbaum, Paul R., and Donald B. Rubin. "The central role of the propensity score in observational studies for causal effects." *Biometrika* 70.1 (1983): 41-55.

Rosenbaum, Paul R., and Donald B. Rubin. "Constructing a control group using multivariate matched sampling methods that incorporate the propensity score." *The American Statistician* 39.1 (1985): 33-38.

Sanbonmatsu, Lisa, et al. "Moving to Opportunity for Fair Housing Demonstration Program--Final Impacts Evaluation." (2011).

Sharkey, Patrick. 2010. "The Acute Effect of Local Homicides on Children's Cognitive Performance." *Proceedings of the National Academy of Sciences* 107:11733-11738.

Spirtes, P., Glymour, C. N., & Schein, R. ([1993] 2001). Causation, prediction, and search (2nd ed.). Cambridge, MA:MIT Press.

Winship C, Elwert F. "Effect Heterogeneity and Bias in Main-Effects-Only Regression Models". In Dechter R, Geffner H, Halpern JY Heuristics, Probability and Causality: A Tribute to Judea Pearl College Publications; 2010.

Winship C, Morgan S. "Bringing Context and Variability Back in to Causal Analysis". In Kincaid H Oxford Handbook of the Philosophy of the Social Sciences Oxford University Press; 2012.

Wood, Robert G., Sheena McConnell, Quinn Moore, Andrew Clarkwest, and JoAnn Hsueh. "Strengthening Unmarried Parent's Relationships: The Early Impacts of Building Strong Families." Princeton, NJ, Mathematica Policy Research, 2010.

Wood, Robert G., Quinn Moore, and Andrew Clarkwest. "BSF's Effects on Couples Who Attended Group Relationship Skills Sessions: A Special Analysis of 15 – Month Data." Princeton, NJ: Mathematica Policy Research, May 2011.

Wooldridge, J. M. 2010. Econometric Analysis of Cross Section and Panel Data . 2nd ed. Cambridge, MA: MIT Press. .

Wooldridge, Jeffrey. *Introductory econometrics: A modern approach*. Cengage Learning, 2012.

Xie, Yu. "Values and limitations of statistical models." *Research in social stratification and mobility* 29.3 (2011): 343-349.

Xie, Yu, Jennie E. Brand, and Ben Jann. "Estimating heterogeneous treatment effects with observational data." *Sociological methodology* 42.1 (2012): 314-347.

Appendix A. Simulation Process

This section explains the mathematical problem solved by simulated agents in the creation of the test data. Agents have preferences over consumption, leisure, and high school completion. They make choices each period over the allocation of time towards, education, labor or leisure given constraints on time availability and the consumption benefits of wage labor. The sequence of choices determines whether they graduate in four periods. This provides a non-linear data generating process with computable counterfactual outcomes, where we can assess our ability to make correct causal inference using the linear models presented above.

Formally, simulated agent i solve.

$$V(l^a_{t,i}, l^w_{t,i}, x_{t,i}, e_{t,i}, c_{t,i}, a_i) = \underset{\{l^a_{t,i},\ l^w_{t,i},\ x_{t,i}\}}{\operatorname{argmax}} \left(x_{t,i}\right)^{\alpha_i}\left(1 - l^a_{t,i} - l^w_{t,i}\right)^{\beta_i} + G_{i,t} \cdot \mathbb{1}\left(l^a_{t,i} \geq \frac{c_{t,i}}{a_i}\right) \tag{4}$$

$$\text{subject to:} \qquad x_{t,i} \leq l^w_{t,i} w_i + e_i \tag{5}$$

$$0 \leq 1 - l^a_{t,i} - l^w_{t,i} \tag{6}$$

$$x_{t,i} \geq 0,\ 0 \leq l^w_{t,i} \leq 1\ , 0 \leq l^a_{t,i} \leq\ 1, \alpha > 0, \beta > 0 \tag{7}$$

The objective function (4) takes the form of a Cobb-Douglas utility function expressing preferences over consumption, $x_{t,i}$, and leisure, $1 - l^a_{t,i} - l^w_{t,i}$, plus an additional additive component capturing gains from graduation. Constraint (5) is a budget constraint limiting present period consumption to what is affordable given endowment e_i and earned income, $l^w_{t,i} w_i$, where w_i is a wage rate.. Endowments are a small but fixed percentage of household income, while wage offers are random. $G_{i,t}$ is the time t value

that agent i places on eventual graduation. The indicator function shows that agent i only expects to receive $G_{i,t}$ if he or she supplies some minimum amount of academic labor. Constraints (6) and (7) stipulate that all time and consumption allocations must be non-negative.

The solution to this problem takes the following form given state $s_{t,i} = \{e_{t,i}, c_{t,i}, a_{i,} G_{i,t}\}$:

$$l_{t,i}^{a*} = \begin{cases} \frac{c_{t,i}}{a_i} & if \quad V\left(l_{t,i}^{a} = 0, l_{t,i|l_{t,i}^{a}=0}^{w*}, x_{t,i|l_{t,i}^{a}=0}^{*}, s_{t,i}\right) \leq V\left(l_{t,i}^{a} = \frac{c_{t,i}}{a_i}, l_{t,i|l_{t,i}^{a}=\frac{c_{t,i}}{a_i}}^{w*}, x_{t,i|l_{t,i}^{a}=\frac{c_{t,i}}{a_i}}^{*}, s_{t,i}\right) and \ \frac{c_{t,i}}{a_i} \leq 1 \\ 0 & otherwise \end{cases} \tag{8}$$

$$l_{t,i}^{w*} = \frac{w_i \alpha_i \left(1 - l_{t,i}^{a*}\right) - e_i \beta_i}{w_i (\beta_i + \alpha_i)} \tag{9}$$

$$x_{t,i}^{*} = \alpha_i \left(\frac{w_i \left(1 - l_{t,i}^{a*}\right) + e_i}{\beta_i + \alpha_i}\right) \tag{10}$$

The intuition behind this solution is as follows. The return function for academic labor is a non-differentiable step function and requires special care for that reason. There are only three possible optimal values for academic labor. First, one could supply $\frac{c_{t,i}}{a_i}$ units of academic labor, which is just enough to receive the expected return G. Any time committed beyond this amount has no return, and would be better spent towards wage labor or leisure since $\alpha_i > 0$ and $\beta_i > 0$. If it turns out that the simulated agent cannot feasibly supply the desired amount of academic labor such that

$$\frac{c_{t,i}}{a_i} > 1$$

or that the agent has a higher present period gain if she devotes the time to wage labor or leisure such that

$$V\left(l_{t,i}^{a} = 0, l_{t,i|l_{t,i}^{a}=0}^{w*}, x_{t,i|l_{t,i}^{a}=0}^{*}, s_{t,i}\right) > V\left(l_{t,i}^{a} = \frac{c_{t,i}}{a_i}, l_{t,i|l_{t,i}^{a}=\frac{c_{t,i}}{a_i}}^{w*}, x_{t,i|l_{t,i}^{a}=\frac{c_{t,i}}{a_i}}^{*}, s_{t,i}\right)$$

then $l_{t,i}^{a*} = 0$ must be the optimal academic labor supply. In this case, any time allocation above 0 has a higher return as wage labor time or leisure time. From here, we utilize the concavity, continuity, and differentiability of the return function in $\{l_{t,i}^{w}, \ x_{t,i}\}$ to solve for (9) and (10) in terms of parameters and $l_{t,i}^{a*}$. These formulas are sufficient for calculating the current period return to continuing education, and the current period return to dropping out. Simulated agents choose the option with the highest present period return.

This model adds interesting dynamics to the data. First, individual actions are highly sensitive to specifications of $\alpha_i, \beta_{i,}$ and $G_{t,i}$ all of which are unobserved by the researcher. Also, these parameters are correlated with familial and community characteristics. Endowments e_i are a function of household income. The dependence of these parameters on family characteristics should lead to an estimable degree of intergenerational transmission of advantage.

Table 1. Time-Varying Treatment Exposure and Outcome Measurement

		Outcome measurement	
		Time-invariant	Time-varying
Treatment exposure	Time-invariant	$Case\ 1$ $\delta_i = y_i^{d=1} - y_i^{d>1}$	$Case2$ $\delta_{i,v} = y_{i,v}^{d=1} - y_{i,v}^{d>1}$
	Time-varying	$Case\ 3$ $\delta_i^{t,T} = y_i^{d=t} - y_i^{d>T}$	$Case\ 4$ $\delta_{iv}^{t,T} = y_{i,v}^{d=t} - y_{i,v}^{d>T}$

Table 2. Descriptive Statistics for Simulation Data

		1980 Cohort Mean/Proportion	2000 Cohort Mean/Proportion
Gender			
	Female	49.65	50.76
	Male	50.35	49.24
Race/Ethnicity			
	White	39.5	39.42
	Black	25.23	25.31
	Asian	10.91	10.9
	Hispanic	24.36	24.38
Mother's Education			
	Less Than High School	39.88	29.55
	High School Degree	36.74	23.71
	College Degree	14.98	26.71
	Graduate/Professional Degree	8.4	20.04
Father's Education			
	Less Than High School	38.71	29.08
	High School Degree	38.23	24.81
	College Degree	15.25	26.18
	Graduate/Professional Degree	7.82	19.94
Household Income			
	Mean	$45,702	$59,709
	<$20,000	21.68	16.52
	$20,000 to $50,000	37.68	32.45
	$50,000 to $100,000	29.91	29.98
	$100,000 to $200,000	10.73	17.2
	$200,000+	0	2.34
Community of Residents			

	Community 1	20.94	26.16
	Community 2	23.46	22.38
	Community 3	19.33	17.04
	Community 4	14.76	12.62
	Community 5	21.51	21.79
Household Structure			
	Both Parents, Married	77.05	69.87
	Mother headed HH	15.79	22.92
	Both Parents, Cohabiting	7.16	7.2
Model Outcomes			
	Periods Employed	4.00	3.95
	Total Wage Labor Over 4 Periods	1.23	1.23
	Total Academic Labor Over 4 Periods	0.79	0.82
	Four Period Graduation Rate	0.56	0.59
	Total Periods Enrolled in School	3.38	3.52
N			
	Respondents	7,003	7,023
	Families	5,250	5,265
	Respondents with siblings	3,290	3,296
	Families with multiple kids	1,537	1,538

Table 3. ETV Probabilities

Brackets indicate periods where a policy was enacted that lead to a sharp change in exposure rates.

1980s ETV Probabilities By Community and Period

		Community				
		1	2	3	4	5
Period	1	3	2	2	4	6
	2	3	1	2	4	6
	3	3	2	[4]	[8]	[12]
	4	3	1	[4]	[8]	[12]

2000s ETV Probabilities by Community and Period

		Community				
		1	2	3	4	5
Period	1	3	2	2	8	12
	2	3	1	2	8	12
	3	3	2	[4]	[4]	[6]
	4	3	1	[4]	[4]	[6]

Table 4a. Data True Effect Calculations, 1980s

	Observed Mean	Counterfactual Mean	ETV Effect	Percent Change Due to Treatment
Total years employed during HS	4.000	4.000	0.000	0.000
Total Hours devoted to wage labor	1.223	1.253	-0.030	-0.024
Total hours devoted to academic labor	0.602	0.749	-0.147	-0.196
Enrollment by Period	0.111	0.802	-0.692	-0.863
Graduated from HS in 4 years	0.321	0.553	-0.232	-0.419

Table 4b. Data True Effect Calculations, 2000s

	Observed Mean	Counterfactual Mean	ETV Effect	Percent Change Due to Treatment
Total years employed during HS	3.914	3.950	-0.036	-0.009
Total Hours devoted to wage labor	1.158	1.257	-0.099	-0.078
Total hours devoted to academic labor	0.769	0.790	-0.021	-0.027
Enrollment by Period	0.768	0.860	-0.093	-0.108
Graduated from HS in 4 years	0.545	0.562	-0.018	-0.031

Table 5a. Counterfactual and Estimated Effects, 1980s

Outcome	Method	True Counterfactual Effect Estimate	Uncontrolled Estimate (Std Error)		Controlled Estimate (Std Error)		True Estimate within 95% C.I.	Estimate has the correct sign
Periods Employed		0						
	OLS Regression		0		0		+	≈
			0		0			
	Propensity Score Matching				0		+	≈
					0			
Total Wage Labor		-0.03						
	OLS Regression		-0.004		-0.010			≈
			(0.0099)		(0.0072)			
	Fixed Effects				0.003			
					(0.0011)			
	Instrumental Variables		0.288	***	0.349	***		
			(0.0244)		(0.0237)			
	Propensity Score Matching		---		-0.047	**	+	≈
			---		(0.0145)			
Total Academic Labor		-0.147						
	OLS Regression		-0.217	***	-0.209	***		≈
			(0.0100)		(0.0085)			
	Fixed Effects		-0.176	***	---			≈
			(0.0029)		---			
	Instrumental Variables		-0.728	***	-0.388	***		≈
			(0.0378)		(0.0405)			
	Propensity Score Matching		---		-0.121	***	+	≈
			---		(0.0171)			

Enrollment by Period	-0.692						
OLS Regression		-0.673	***	-0.680	***	+	≈
		(0.0067)		(0.0066)			
Fixed Effects		-0.653	***	---			
		(0.0068)		---			
Instrumental Variables		-3.578	***	-2.693	***		≈
		(0.1625)		(0.1869)			
Propensity Score Matching		---		-0.733	***		≈
		---		(0.0152)			
Graduated in four years	-0.232						
OLS Regression		-0.295		-0.286	***		≈
		(0.0159)		(0.0154)			
Probit		-0.762		-0.637	***		≈
		(0.0429)		(0.0462)			
Logit		-1.227		-1.043	***		≈
		(0.0705)		(0.0763)			
Instrumental Variables		-2.147		0.000		+	
		(0.1536)		(0.1797)			
Propensity Score Matched		---		-0.188	***	+	≈
		---		(0.0236)			

Notes: Models include the following control variables: Mother's education, father's education, race, family structure, household income, community of residence, and simulated respondent wage offers. Asterisks and daggers for p-value significance have the standard meaning: p<.001 = ***, p<0.01=**, p<0.05=*, p<0.1=†

Table 5b. Counterfactual and Estimated Effects, 2000s

Outcome	Method	True Counterfactual Effect Estimate	Uncontrolled Estimate (Std Error)		Controlled Estimate (Std Error)		True Estimate within 95% C.I.	Estimate has the correct sign
Periods Employed		-0.036						
	OLS Regression		-0.038	**	-0.042	**	+	≈
			(0.0147)		(0.0138)			
	Propensity Score Matching				-0.059	**	+	≈
					(0.0211)			
Total Wage Labor		-0.099						
	OLS Regression		-0.085	***	-0.086	***		≈
			(0.0123)		(0.0078)			
	Fixed Effects		---		-0.012	***		≈
			---		(0.0010)			
	Instrumental Variables		0.324	*	-0.078	**	+	≈
			(0.1324)		(0.0245)			
	Propensity Score Matching		---		-0.106	***		≈
			---		(0.0179)			
Total Academic Labor		-0.021						
	OLS Regression		-0.058	***	-0.061	***		≈
			(0.0095)		(0.0081)			
	Fixed Effects		---		-0.02645392	***	+	≈
			---		0.00288027			
	Instrumental Variables		-1.406	***	-0.006		+	≈
			(0.3281)		(0.0359)			
	Propensity Score Matching		---		-0.036	*	+	≈
			---		(0.0147)			
Enrolment by Period		-0.093						
	OLS Regression		-0.119	***	-0.121	***		≈
			(0.0092)		(0.0078)			
	Fixed Effects		---		-0.081	***		≈
			---		(0.0061)			
	Instrumental Variables		-7.512	***	0.525	***		

		(1.6940)		(0.1096)			
Propensity Score Matching		---		-0.109	***	+	≈
		---		(0.0158)			
Graduated in four years	-0.018						
OLS Regression		-0.057	***	-0.057	***		≈
		(0.0158)		(0.0152)			
Probit		-0.146	***	-0.071			≈
		(0.0404)		(0.0435)			
Logit		-0.235	***	-0.114			≈
		(0.0647)		(0.0708)			
Instrumental Variable		---		0.000			
		---		(0.1736)			
Propensity Score Matching		---		-0.045	†		≈
		---		(0.0231)			

Notes: Models include the following control variables: Mother's education, father's education, race, family structure, household income, community of residence, and simulated respondent wage offers. Asterisks and daggers for p-value significance have the standard meaning: p<.001 = ***, p<0.01=**, p<0.05=*, p<0.1=†

Figure 1. DAG Causal Model of Effects of ETV on High School Graduation

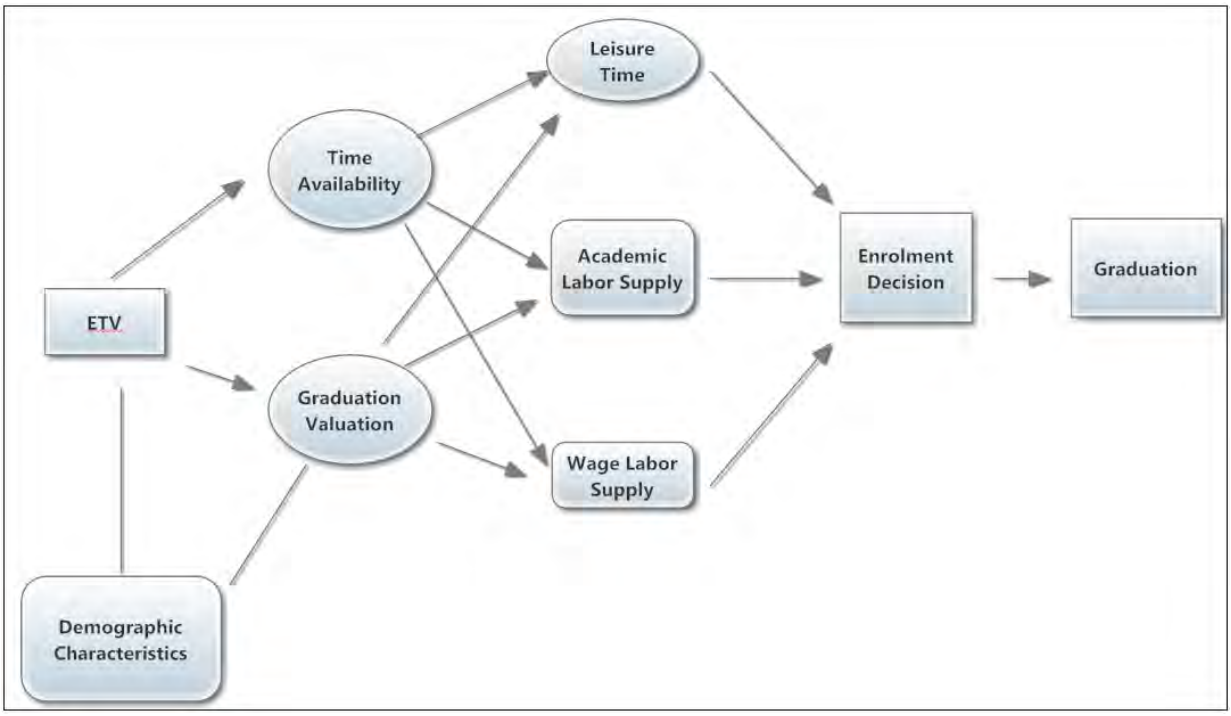

Figure 2. Proportion of Sample Experiencing ETV by year