CHAPTER 19

Statistical Models of Life Events and Criminal Behavior*

D. WAYNE OSGOOD

The goal of developmental and life course criminology is to understand patterns of crime and delinquency over the life course. To date, research in this field has devoted a great deal of attention to describing patterns of change in the dependent variable over different ages, often in the form of trajectories or growth curves of offending in relation to age (LeBlanc and Loeber 1998; Piquero et al. 2007). Closely tied to such studies is a sizable body of research investigating potential predictors of differences in trajectories (e.g., Nagin et al. 1995; Nagin and Tremblay 1999, 2005).

Focusing on growth curves or trajectories has often been associated with an emphasis on the role of early experience and personality traits in shaping the course of development of crime. According to Laub and Sampson (2003), however, criminologists need to balance such work with greater attention to the connections between crime and later events in people's lives. Indeed, the longitudinal data used to study trajectories is equally suitable for this purpose as well. Furthermore, both traditional and developmental/life course theories of crime hypothesize a wide variety of effects on crime from life events, which recent studies have begun to test (for a review, see Siennick and Osgood 2007). For instance, both traditional (Hirschi 1969) and age-graded (Sampson and Laub 1993) versions of social control theory predict that changes in social bonds will affect offending (King et al. 2007; Sampson et al. 2006), social learning theories (Akers 1977; Sutherland and Cressey 1955) predict that offending will increase after a switch to a more delinquent peer group (Warr 1993), and generalized strain theory (Agnew 1992) predicts that experiencing stressful events will promote offending (Slocum et al. 2005).

The aim of this chapter is to facilitate research on life events and crime by presenting a set of statistical tools for analyzing the relationship between events in people's lives and changes in their levels of offending. The presentation will concentrate on events that constitute a categorical change in one's life, such as marriage, gaining employment, or entering a treatment or service program. For the most part, however, these methods apply equally well to

* A version of this chapter was presented at The Social Contexts of Pathways in Crime: Methods and Analytical Techniques, The Second Annual SCoPiC Conference, Cambridge, UK, June 2005.

A.R. Piquero and D. Weisburd (eds.), *Handbook of Quantitative Criminology*,
DOI 10.1007/978-0-387-77650-7_19,

studying changes of degree in more continuous variables such as marital commitment or job satisfaction, which are also central to the study of crime over the life course (e.g., Sampson and Laub 1993).

The statistical approach I present is meant for analyzing longitudinal panel data with repeated measures of both crime and a time-varying explanatory variable reflecting the event of interest (e.g., not married versus married). This research design is typically "observational" in the sense that it captures only naturally occurring variation in the explanatory variable, in contrast to a study in which the variation comes from random assignment or is affected by a strictly exogenous source, such as a policy change. Though the techniques presented here are also useful for longitudinal studies with random assignment (Esbensen et al. 2001; Osgood and Smith 1995), I will give special attention to issues of causal inference that arise with observation data. Observational research designs can never yield definitive proof of causality, but the tools presented here offer means of ruling out several important types of competing explanations and thereby strengthening the plausibility of a causal interpretation.

AN INITIAL MODEL OF THE EFFECTS OF EVENTS

The key feature that distinguishes these statistical models is the *time varying covariate*, which is simply an explanatory variable that can vary over time for a person. The idea of a time varying covariate implies a longitudinal research design that follows a sample of individuals through multiple measurements over some period. We study the effects of events by analyzing how the outcome of interest, crime, relates to change over time in a variable reflecting the event of interest, such as whether or not the respondent is married.

The most basic model for accomplishing this is simply:

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + e_{ti} \tag{19.1}$$

Equation (19.1) looks very much like a standard bivariate regression equation, but there is one subtle difference. In the standard version, variables have just one subscript, but here they have two, i and t. The additional subscript means that we are differentiating observations not just in terms of the person being studied (i), but also in terms of the occasion or time of measurement (t). Thus, (19.1) specifies that both the outcome variable, Y, and the explanatory variable, X, are measured for each wave of data that a person contributes to the analysis, and that their values can vary within individuals over time.

Straightforward interpretations for the regression coefficients of (19.1) follow from coding the explanatory variable as a dummy variable equaling 0 before the event occurs (e.g., unmarried) and 1 after it has occurred (e.g., married). In this case, β_0 will reflect the mean of the outcome variable before the event (the crime rate when respondents are not married), and β_1 will capture its difference from the mean of the outcome after the event. Thus, a large negative value for β_1 would indicate considerably less crime among married respondents than among unmarried respondents.

STATISTICAL CONCERNS

The focus of this chapter is on variations of (19.1) that flexibly capture potential patterns of change associated with events and that rule out important alternative interpretations of the relationships between events and outcomes. Before turning to these matters, it is important to consider some technical issues that mean (19.1) must be estimated using specialized statistical techniques. Because these issues are not the focus of this chapter, this section only briefly explains the main concerns and their standard solutions, referring the reader to other sources for more thorough discussions.

Basic statistical models, such as ordinary least squares regression (OLS), assume independence among observations. For (19.1), this assumption implies that there should be no systematic relationships among the residuals, e_{ti}, for different observations, so that knowing the value e for any one observation should be of no help for predicting its value for any others. Data from longitudinal panel studies almost always violate this assumption in two ways. First, even modest stability in individual behavior means that people are generally more similar to themselves across different occasions than they are to most other people. Second, individual change over time tends to be at least somewhat gradual rather than totally haphazard, with the result that observations from the same individual that are closer to each other in time are, on average, more similar than are observations that are farther apart. The residuals, e_{ti}, would be independent only if a regression model fully accounted for these response patterns, and it would be inappropriate to assume in advance that it did. Thus, analyzing longitudinal panel data requires a statistical model that takes into account dependence among observations that arises from both individual differences in average response levels and serial correlation.

One can allow for dependence among observations due to individual differences in average response level by dividing the original residual term, e_{ti}, into two components, one capturing that consistency across time for each person, u_i, and the other the time-specific variation around that average level, r_{ti}. It is then plausible to assume that the u_i residuals are independent across individuals and that the r_{it} residuals are independent within each individual over time (ignoring serial correlation for the moment). This division of e_{ti} changes (19.1) to (19.2):

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + u_i + r_{ti} \tag{19.2}$$

An alternative way of representing consistent individual differences that are not accounted for in the model is to make the intercept term unique to each person by adding the subscript i, yielding β_{0i}. Equation (19.3) illustrates this conception of the problem:

$$Y_{ti} = \beta_{0i} + \beta_1 X_{ti} + r_{ti} \tag{19.3}$$

Equation (19.4) shows that the two approaches are mathematically equivalent, and thus the difference between (19.2) and (19.3) is a matter of notation rather than substance.

$$\beta_{0i} = \beta_0 + u_i \tag{19.4}$$

There are two widely available methods for estimating this statistical model. The fixed effects[1] approach directly estimates a separate intercept, β_{0i}, for each individual. This method, in effect, adds to (19.1) a separate dummy variable for each person, resolving the dependence among repeated observations by fully removing all individual differences in average response

[1] The name "fixed effects" is unfortunate because it is used for many other purposes in statistics as well. For instance, it should not be confused with the use of the term "fixed effect" to refer to the regression coefficients estimated for explanatory variables in random effects models (as in printout from the HLM program).

levels. The alternative is the random effects approach, which instead estimates the variance of the residuals u_i and uses that estimate to adjust the regression coefficients and their standard errors.

Both of the methods are quite useful, and each has advantages and disadvantages. Overall, the fixed effects method requires fewer assumptions, while the random effects method provides greater statistical power and is applicable to a broader range of problems. For more detailed comparisons of the two approaches, see Allison (2005), Johnson (1995), and Peterson (1993). Both methods can be found in advanced statistical software such as STATA and LIMDEP, and the random effects model can be estimated through multilevel regression programs such as MLwiN, HLM, and SAS PROC MIXED.

Even after addressing dependence due to individual differences in average responses, some dependence is likely to remain among the residuals r_{ti} because observations that are closer in time will be more similar than those that are more widely separated. There are several ways to tackle this serial correlation. The typical growth curve modeling approach is to expand the random effects model just described with additional residual variance terms that allow for variation across individuals in the relationship of age or time to the outcome (Bryk and Raudenbush 1987). This more elaborate version is referred to as a random coefficient model as opposed to a random intercept model. Following the tradition of time series research, one could instead directly model the serial dependence, and the most common means of doing so is by estimating the autocorrelation between the residual terms of adjacent time points. A final option is simply to estimate and adjust for all of the across-wave correlations of the residuals. Though this last approach is guaranteed to fit the data well, it is less useful when there are many observations per person or the timing of observations is highly variable. All three approaches for addressing serially correlated error have been implemented for both multilevel regression models (Raudenbush 2001b) and structural equation models for latent growth curves (Curran and Bollen 2001; Rovine and Molenaar 2001). Raudenbush's (2001b) discussion of serially correlated error in analyses of longitudinal panel data will prove useful for readers who would like to know more about the topic.

Measures of crime and deviance usually have highly skewed and discrete distributions that violate other assumptions of standard statistical models such as homogeneity of residual variance and consistency between fitted values and actual means (Osgood and Rowe 1994). For cross-sectional data, generalized linear models, such as logistic and Poisson regression, will usually resolve these problems. These generalized models are also available within multilevel regression (e.g., Raudenbush and Bryk 2002; Snijders and Bosker 1999), providing a flexible framework for analyzing longitudinal panel data on crime and deviance that is suitable for all of the models presented in this chapter. Readers seeking additional background on generalized linear models should see sources such as Agresti (2007) or Long (1997).

In sum, the statistical issues that arise from a longitudinal research design and from the nature of measures of crime and deviance require the use of specialized statistical methods briefly discussed here. The remainder of the chapter omits these elements from the models presented, however, to reduce complexity not directly relevant to issues being discussed.[2]

[2] For instance, I will present models with the standard residual term e_{ti}, though the reader should assume that these terms are not independent, and that in actual data analysis, they would be replaced by more complex composite residuals, such as $e_{ti} = u_{0i} + u_{1i}$ Age $+ r_{ti}$.

FOCUSING ON CHANGE

In (19.1), our initial model, the regression coefficient β_1 indexes the association of the event variable with the outcome of interest. For the example of marital status and crime, this coefficient will contrast the mean of the crime measure for all observations of married people with its mean for all observations of unmarried people. Unfortunately, the difference between these means will reflect not only the impact of the event of interest, but also any preexisting differences in offense rates that are associated with the event variable. For instance, a person who never marries contributes only to the mean for the unmarried status, while a person who was married for the entire study contributes only to the mean of the married status. Thus, much like a cross-sectional analysis, β_1 from (19.1) is an undifferentiated amalgam of within-person change associated with the event and preexisting differences between people who do and do not experience the event.

If we are interested in the impact of an event, then we would like to eliminate the contribution of prior differences and focus instead on the association between within-individual change on the event variable and within-individual change in offending (Horney et al. 1995). One way to accomplish this would be to reformulate our regression equation as follows:

$$\left(Y_{ti} - \overline{Y}_{\bullet i}\right) = \beta_1 \left(X_{ti} - \overline{X}_{\bullet i}\right) + e_{ti} \tag{19.5}$$

Equation (19.5) eliminates all stable individual differences on both the event variable and the outcome by subtracting each individual's mean across time from his or her scores for both X and Y from each wave of data.[3] Thus, β_1 will reflect only the within-individual association between X and Y, contrasting each person's offense rate when married with his or her offense rate when unmarried, and pooling that information across the sample. People who never marry or who are always married will not contribute to β_1 because, for them, the transformed explanatory variable, $(X_{ti} - \overline{X}_{\bullet i})$, has a constant value of zero.

There are several straightforward ways to obtain estimates of within-individual relationships that are equivalent to (19.5). The first is through the fixed effects approach to addressing dependence due to stable individual differences, discussed above. The dummy variables for every person in the analysis fully account for all individual differences in average response, restricting the estimates of relationships for time-varying explanatory variables to within-individual change over time. For examples of criminological studies using this approach to obtaining within-person estimates, see Osgood et al. (1996) and Paternoster et al. (2003). Another approach is to subtract the individual means of the explanatory variable (also known as group mean centering), as was done by Horney et al. (1995):

$$Y_{ti} = \beta_0 + \beta_1 \left(X_{ti} - \overline{X}_{\bullet i}\right) + e_{ti} \tag{19.6}$$

In this case, the explanatory variable has a mean of zero for each person, so it cannot explain differences between people, but rather only within-individual variation in the outcome. A third approach is to add the individual mean on X as an additional explanatory variable:

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + \beta_2 \overline{X}_{\bullet i} + e_{ti} \tag{19.7}$$

[3] Note that (19.5) has no constant term, β_0. None is needed because subtracting the individual means constrains both the left- and right-hand sides of the equation to have means of zero.

This approach capitalizes on a basic principle of regression analysis that the coefficient for each variable is determined entirely by the portion of its variance that is independent from the other variables in the model. The variance in X that is independent of individual means over time would be within-individual change in X. These three strategies will produce logically equivalent and (with reasonably large samples) nearly identical estimates of the effects of the time-varying explanatory variable (Allison 2005; Bushway et al. 1999; Raudenbush and Bryk 2002).

Taking these simple steps to focus the analysis on within-individual change over time has enormous methodological importance. As Allison (1990) has shown, this strategy is a much more effective means of adjusting for prior differences than is treating earlier measures of the outcome as covariates. Analysis of within-individual change totally eliminates the possibility that any stable individual characteristic can account for the estimated effect of the time-varying explanatory variable, without measuring that characteristic and including it in the analysis. No stable factor such as gender, IQ, or self-control can possibly account for a lower rate of crime when people are married compared to their own rates of crime when they were not married.[4] Thus, this approach precludes a broad class of potential selection effects, which is why methodologists view it as one of the most powerful tools for studying causal processes using nonexperimental data (Allison 2005; Greene 2000; Winship and Morgan 1999).[5]

It may appear that controlling for unmeasured variables in this fashion gains us something for nothing, but that is not the case. Rather, this strategy is only possible because of strength inherent in the longitudinal panel research design, which enables us to use "subjects as their own controls," in the old terminology of experimental design. The limits of this design mean that two key explanations remain as competing alternatives to a causal interpretation of the estimates from (19.5) to (19.7). First, these within-person estimates control only for stable unmeasured characteristics, but not for other time-varying variables. Thus, these estimates are potentially subject to omitted variable bias due to any time-varying variables omitted from the model that have a causal impact on the outcome and are correlated over time with the variable of interest. For instance, an apparent effect of marriage might be partly or entirely due to an uncontrolled variable such as parenthood or income. Second, this within-person analysis does not preclude the possibility that some or all of the estimated effect of the event on crime is actually due to an influence of crime on the event, which would be a form of simultaneity bias.

MATURATION AND EFFECTS OF EVENTS

The next issue we need to consider in assessing effects of events on crime is the threat to validity that Campbell and Stanley (1966) labeled maturation. Maturation threatens validity when the apparent effect of an event may be due merely to similar age trends for both the event and the outcome. As Raudenbush (2001a: 523) pointed out, longitudinal analyses must take into account that people are naturally growing or changing apart from the effects of the

[4] Of course this reasoning only applies to the extent that a variable really is stable over time. Thus, this approach would not control for effects of changes in self control, if self control varied meaningfully over time, contrary to Gottfredson and Hirschi's claim (Hay and Forrest, 2006).

[5] As Allison (1990) pointed out, true experiments with random assignment are an exception for which this approach is not optimal. In that case, prior differences are attributable to chance rather than genuine group differences, and limiting the analysis to within-individual change unnecessarily sacrifices statistical power.

variables of interest. For instance, because rates of marriage increase with age from the late teens into middle adulthood while rates of crime decrease, there may be a negative association between the two, even if marriage has no impact on crime. Similarly, parental supervision could be associated with delinquency from preadolescence through middle adolescence only because the former decreases while the latter increases over that period. Though such similarities in age trends no doubt spark life course scholars' interest in potential explanatory variables, they also mean that we must construct our analyses to insure that estimated effects of those variables are not merely an artifact of that similarity.

Age is, of course, a time-varying variable, and thus maturation constitutes a potential omitted variable bias for analyses of the effects of events. As with other relevant but omitted variables, the straightforward way to address maturation as an alternative explanation is to incorporate the age trend in the regression model, such as:

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + \beta_2 \text{Age}_{ti} + \beta_3 \text{Age}_{ti}^2 + \beta_4 \overline{X}_{\bullet i} + e_{ti} \tag{19.8}$$

The two new terms in (19.8) allow for a quadratic age trend, thereby combining a growth curve model with (19.7). The quadratic trend will be appropriate for some problems, too complex for some, and too simple for others. Because developmental trends are rarely linear over extended periods, it is important to be sure that the form of the age trend in the model is well suited to the data, as reflected in a close match between the fitted age trend and age-specific means. Several sources provide useful guidance for specifying this aspect of the model (e.g., McClendon 1995; Raudenbush and Bryk 2002; Singer and Willett 2003).

Adding age trends to the analysis also provides a vehicle for assessing whether an event accounts for some, or all, of the age trend in offending, which is one of the most central topics of the study of crime and the life course (Osgood 2005). Hirschi and Gottfredson's (1983) influential paper on age and crime claimed that social variables cannot account for the dramatic changes in rates of offending over the life course. Since that time, developmental and life course criminologists have offered many theories that attempt to do so (e.g., Moffitt 1993; Sampson and Laub 1993; Thornberry and Krohn 2005), but there have been surprisingly few empirical tests of the success of those theories in this regard (but see Osgood et al. 1996; Warr 1993).

To determine how well the event variable, X, accounts for age trends in offending, results for (19.8) can be compared to those for a reduced form model that excludes that variable:

$$Y_{ti} = \beta_0 + \beta_2 \text{Age}_{ti} + \beta_3 \text{Age}_{ti}^2 + e_{ti} \tag{19.9}$$

An event such as marriage will account for the average age trend in offending to the degree that it is strongly associated with offending and that its age trend matches the age trend in offending (Hirschi and Gottfredson 1985). This will be reflected in reductions of the coefficients for Age and Age^2 between (19.8) and (19.9). Because the age trend is curvilinear, the extent of mediation is more easily seen in a graph such as Fig. 19.1 than in the coefficients themselves. The more that the explanatory variable accounts for the age trend, the flatter the age trend after adjusting for that variable, in comparison to the original or overall trend. A useful basis for judging the proportion of the age trend that has been explained is to compute fitted values based on age from (19.8) and (19.9) (by applying the coefficients for age to the values of age in the dataset) and compare the resulting standard deviations (Osgood et al. 1996).

Time-varying explanatory variables such as marriage, employment, or peer relations not only have the potential to explain the average age trend in offending, but also to account for

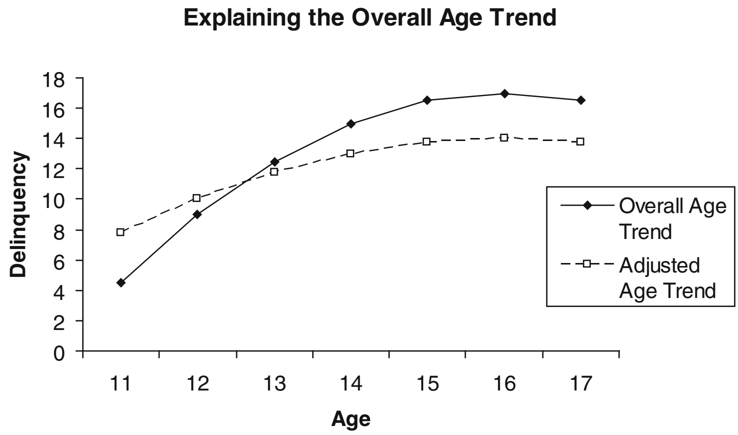

FIGURE 19.1. Illustration of explaining a portion of the age trend by controlling for a time varying covariate.

individual differences in patterns of change over time, as they are reflected in individual trajectories or growth curves. In growth curve models, those individual differences are expressed in the variance components for the polynomial terms for age. Accordingly, an event such as becoming a parent succeeds in explaining how people differ in their offending trajectories to the degree that adding it to the model reduces those variance components (i.e., (19.9) versus (19.8)). For an illustration, see Jacobs and colleagues' (2002) analyses of self-concepts and values concerning achievement.

Studying More Complex Patterns of Change

The models discussed so far assume that the effect of an event takes a specific, simple form. This section of the chapter explains several ways to expand the basic model in order to capture more complex patterns of change. As before, some of these variations will also prove useful for ruling out alternative explanations, thereby increasing the plausibility of viewing results as reflecting a causal effect of the variable of interest.

The preceding regression equations index the effect of an event as a mean difference between observations before the event and those after it. When applied to metric (rather than dichotomous) time-varying explanatory variables, the results will reflect the difference in means between data points that differ by one unit on that variable. As such, these models imply that the effect of the event or change on X is immediate, that it is constant over age and time, and that it applies equally to different subgroups. Figure 19.2 presents hypothetical data to illustrate the assumed pattern for two individuals, one with a higher initial offense rate who experiences the event of interest at age 21 and the other with a lower initial offense rate who experiences it at age 26. The solid lines represent the patterns of change over time accounted for by a model that focuses on within-individual change and adjusts for the overall age trend. The lines dip when the event occurs, indicating that it is associated with a decline in offending. The dashed lines show the pattern of change due to maturation that would be expected if the event did not occur. The difference between dashed and solid lines reflects the impact of the event, which corresponds to β_1 in the above regression equations. As required by these

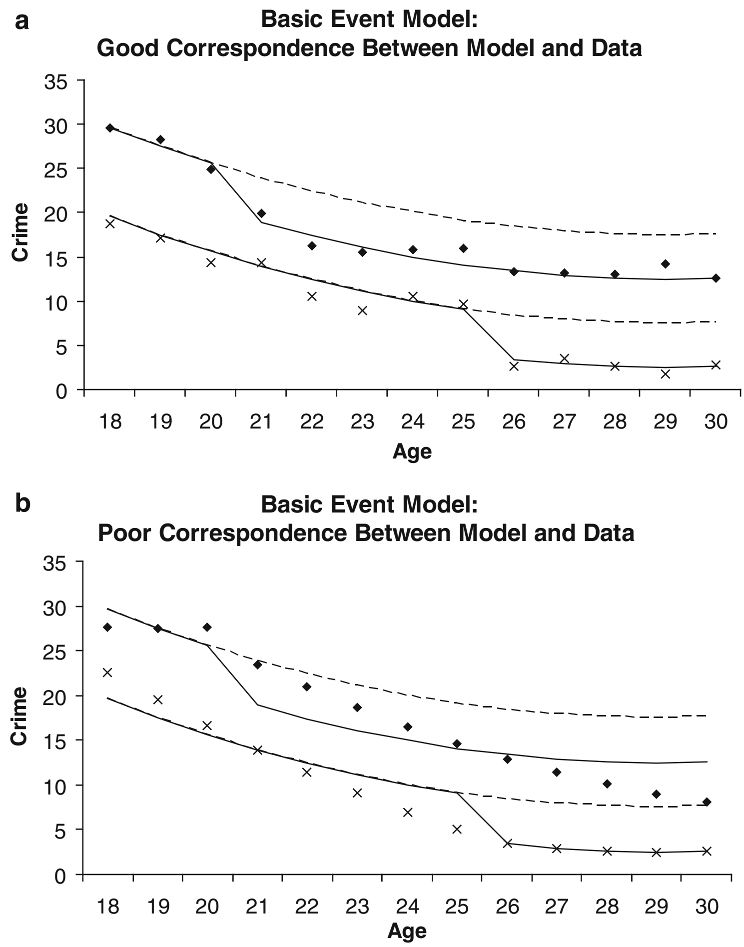

FIGURE 19.2. The form of change inherent in the basic event model, (**a**) good correspondence, (**b**) poor correspondence.

regression models, this difference is immediate, permanent (as long as X stays at 1 rather than 0), and applies equally to both individuals.

The separate points shown in Fig. 19.2 represent the hypothetical raw data, which would be the amount of offending observed for each person, each year. Figure 19.2a illustrates a good correspondence between the model and the data. The data points fall close to the solid lines, and it would be hard to discern any coherent pattern in the discrepancies between the data points and the fitted line from the model. In Fig. 19.2b, however, the discrepancies between the model and the data are systematic and suggest coherent patterns of change inconsistent with the constraints of the statistical model. These systematic departures illustrate some of the limitations of the basic model. For the person with the higher crime rate, the raw data suggest that the event's impact is not immediate and constant, but rather is initially small and grows over time. The effect of the event appears illusory for the person with the lower rate of crime. His or her offending after the event is consistent with the trend in offending before the event, not a departure from it.

Before turning to ways of modifying our basic model in order to capture patterns such as these, it is worth considering advantages that will be lost in more complex models. Because this initial model uses a single parameter to summarize the time-varying covariate's impact across several time points, it has greater statistical power than most other alternatives. It also expresses the relationship in a simple form that is easy to explain, namely, the mean change in the outcome upon the occurrence of the event (or per unit of change in a metric variable such as hours per week spent with spouse). Furthermore, the model's data requirements are so minimal that it is applicable to the simplest of longitudinal studies, even those with only two waves of data. More complex alternatives will require additional data.

AGE VARYING EFFECTS OF EVENTS

First, consider the possibility that the effect of an event might vary with age, which would be inconsistent with the constraints illustrated in Fig. 19.2. Age differences of this sort are a core feature of Thornberry's interactional theory (1987; Thornberry and Krohn 2005), which specifies that variables such as parental attachment and peer delinquency have a greater impact during the developmental periods in which they are most prominent. Indeed, age differences in influences on behavior are central concerns of developmental and life course perspectives.

For regression models of the sort considered here, age variation in effects corresponds to interactions between age and the time-varying explanatory variable. A basic form of this model would be:

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + \beta_2 \text{Age}_{ti} + \beta_3 X_{ti} \times \text{Age}_{ti} + \beta_4 \overline{X}_{\bullet i} + \beta_5 \overline{\text{Age}_{\bullet i}} + \beta_6 \overline{(X_{ti} \times \text{Age}_{ti})}_{\bullet i} + e_{ti} \qquad (19.10)$$

The new coefficient β_3 will reflect change in the effect of X per unit of age. Jang (1999) used models in this form to test interactional theory's main hypotheses about age varying effects. This specific interaction term constrains the effect of X to change by a constant value each year, which may or may not be a good match to the data. One can capture other patterns through interaction terms with age coded in other ways, such as a polynomial or a categorical classification (e.g., 14–18 versus 19–22).

As discussed above, estimates of the effects of events should be restricted to within-individual change in order to reduce the possibility of spurious results due to selection effects. Equation (19.10) accomplishes this by also including terms for the individual means of all time-varying variables, including the interaction term. The alternative strategy of subtracting individual means from the event variables (as in (19.6)) would not require this second version of the interaction term, and thus it is somewhat simpler in this regard. To simplify the presentation in the rest of this chapter, the remaining equations will omit the extra elements needed to limit the analysis to within-individual change.

Figure 19.3 illustrates the consequences of allowing a time varying effect in (19.10). Figure 19.3a shows an effect of parental supervision that is constant across ages (as in (19.8)), which is reflected in the unchanging difference between the higher rate of delinquency for adolescents who receive little supervision in comparison to the lower rate for adolescents who are highly supervised. In Fig. 19.3b, adding the interaction between parental supervision and age allows the effect of supervision to progress from quite sizable at age 11 to minimal at age 19. Figure 19.3b would be consistent with the impact of parental supervision declining with age, perhaps due to norms that parents should grant increasing autonomy to their children as they near the age of home-leaving.

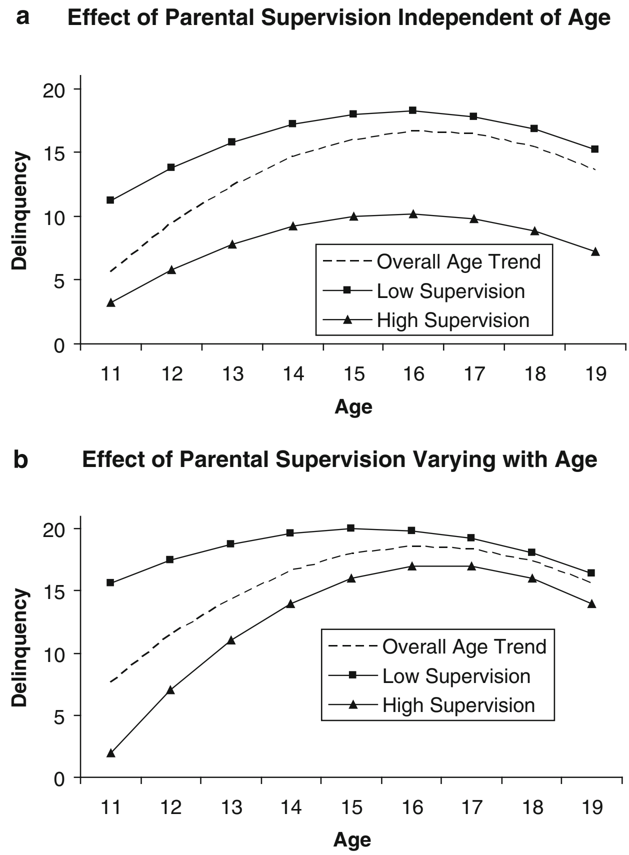

FIGURE 19.3. Illustration of time varying effects of an event, (**a**) effect independent of age, (**b**) effect varying with age.

Figure 19.3 also reflects the relationship between a time-varying explanatory variable and the overall age trend, which appears as a dashed line. The overall trend always falls between the trends for adolescents receiving the two levels of supervision. At younger ages, the overall trend is closer to the line for more highly supervised youth because high supervision predominates. Most parents want to keep close track of where their 11 year old children go and what they do. By age 19, however, few youth are so closely supervised, and accordingly, the average rate of supervision is quite close to the other line. The degree to which parental supervision can explain the age trend is readily apparent in Fig. 19.3a from the steeper slope of the overall age trend relative to the age trends given specific levels of supervision. How well supervision explains the age trend is less apparent from Fig. 19.3b because the trend differs with the level of supervision.

THE TIMING OF AN EVENT'S EFFECTS

Many interesting research questions concern patterns of change inconsistent with the basic model's assumption that an event's effect is immediate and constant over time. For instance, we might hypothesize that the impact of family poverty grows with the length of time the family is poor, that the benefits of a marriage or job emerge gradually (Laub and Sampson 2003; Sampson and Laub 1993), or that the beneficial effect of a treatment program grows during participation in the program and decays after finishing (Osgood and Smith 1995).

To address such possibilities, our model must incorporate information about the timing of an observation relative to the event of interest.

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + \beta_2 \text{Time}(X)_{ti} + \beta_3 \text{Age}_{ti} + e_{ti} \tag{19.11}$$

Equation (19.11) does so by adding a variable, $\text{Time}(X)_{ti}$, that indexes the time elapsed since the event occurred. $\text{Time}(X)_{ti}$ would equal zero from the start of the study until the observation at which the event occurred. It would change to one at the subsequent time point, two at the next time point, three at the time point after that, and so forth. With the data in this form, β_1 will now reflect the average instantaneous change upon the occurrence of the event and β_2 will indicate the amount of increase or decrease in that change per unit of time (relative to the baseline age trend). Another way of putting this is that β_1 is a change in level associated with the event, and β_2 is a change in the slope over time. When the event occurs, the intercept changes from β_0 to $\beta_0 + \beta_1$ and the slope for age changes from β_3 to $\beta_2 + \beta_3$. Osgood and Smith (1995) explain the application of this general modeling strategy to topics in program evaluation.

Examples of the patterns of change possible in (19.11) appear in Fig. 19.4. First, employment brings about an immediate reduction in crime that stays constant over the years in Fig. 19.4a. This pattern corresponds to a change in level (β_1) of -4 combined with 0 change in slope (β_2), in which case (19.11) reduces to the more restrictive basic model of (19.8). In contrast, in the example of Fig. 19.4b, employment brings no immediate crime reduction ($\beta_1 = +1$) but instead carries a gradually increasing benefit ($\beta_2 = -1.2$) that eventually results in a very low crime rate. This pattern could reflect processes such as growing social control from increasing commitment to the job or gradual socialization toward conventional values from time spent with fellow employees who are more prosocial. Finally, in Fig. 19.4c, employment has a large immediate beneficial impact on offending ($\beta_1 = -8$), but that benefit decays over time ($\beta_2 = 1.2$) as the crime rate gradually returns to the baseline.

TIMING OF EFFECTS AND CAUSAL INTERPRETATION. The timing of effects is also relevant to determining whether or not findings are consistent with a causal interpretation. If the change in offending precedes the event, then it is hard to argue that the event caused that change. In contrast, a lag between the event and change in the outcome is often seen as evidence that the effect is causal. Yet, the relevance of such a lag depends on the processes thought to explain offending. For instance, the socialization and accumulation of experience prominent in social learning theory should take time, which would be consistent with delayed effects. We could incorporate this delay in our statistical model if the theory made clear just how long the delay should be, but theories in social and behavioral science rarely have this level of precision. Furthermore, many theories imply relatively instantaneous effects, such as the situational explanation of routine activity theory (Felson 2002; Osgood et al. 1996) or social control theory's emphasis on current bonds to conventional people and institutions (Hirschi 1969). In these cases, a delay between a change in the explanatory variable and a change in the outcome would actually weigh against the theory. Unfortunately, it is especially difficult to distinguish an instantaneous effect of an event on crime from a reverse effect of crime on the event.

FIGURE 19.4. Effects of events on both levels and slopes, (**a**) constant effect, (**b**) growing impact, (**c**) decaying impact.

Regardless of the theory, however, a closer examination of the timing of the relationship between changes in the explanatory variable and changes in offending would be useful for determining whether the data are consistent with viewing the event as a cause of crime. One simple way of doing this is to compare results across three analyses: (a) the first based on the current value of the explanatory variable, X_t, in order to capture the simultaneous relationship (as in (19.8)), (b) the second substituting its earlier value, X_{t-1}, to capture the lagged or delayed relationship, and (c) the third using its later value, X_{t+1}, so that the outcome measure precedes the event. Figure 19.5 shows these three possibilities. A stronger relationship of

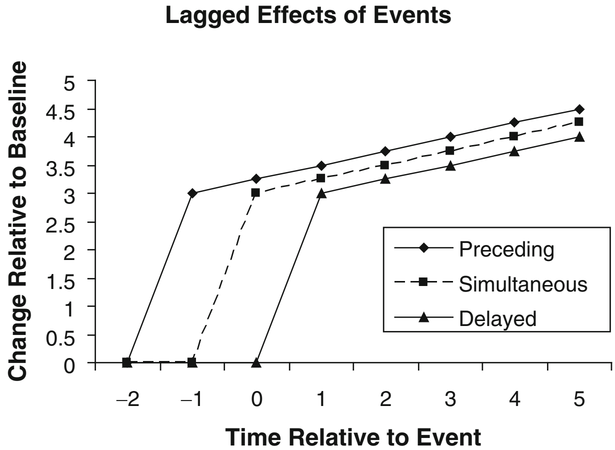

FIGURE 19.5. Lagged effects of events.

offending to the subsequent value of X suggests that the "effect" precedes its supposed cause, so the relationship is likely to be spurious rather than causal.[6]

This strategy is simple, applies to both events and metric (i.e., nondichotomous) time-varying explanatory variables, and is suitable to all of the models discussed so far. It does not provide a very precise picture of the timing of changes in the outcome, however, so it can easily yield ambiguous findings in which there are minimal differences among the three specifications of the timing of X.

For an explanatory variable reflecting an event, we can gain a more precise understanding of the timing of changes in the outcome by following Laub et al.'s (1998) strategy of further reducing constraints on the form of the effect over time. Equation (19.8) assumed the effect was instantaneous and constant over time by representing the event through a single dummy variable indicating its presence or absence. Equation (19.11) allowed the effect to evolve by adding an index of time since the event occurred, but this model still limits the effect to growing or declining in a uniform progression. Laub and colleagues' (1998) approach removes that constraint by representing time in relation to the event through a series of dummy variables, as illustrated in (19.12):

$$Y_{ti} = \beta_0 + \sum_{j=-2}^{5} \alpha_j T_{jti} + \beta_2 \text{Age}_{ti} + e_{ti} \tag{19.12}$$

The dummy variables T_{jit} indicate the timing of each observation t relative to the occurrence of the event for individual i. The index j calibrates this timing, so the variable T_{-1ti} is zero unless this observation is one time unit before the event, in which case, it equals one, T_{2ti} indicates whether this observation is two time units after the event in the same fashion, and so forth. Accordingly, the coefficients α_j will reflect the within-person change (relative to the baseline age trend) for various time points before and after the event.

[6] This reasoning would not apply to cases where it is plausible that a causal effect would precede the measured event. For instance, for marriage this pattern might be interpreted as indicating that the important event is the onset of a committed romantic relationship rather than its culmination in marriage.

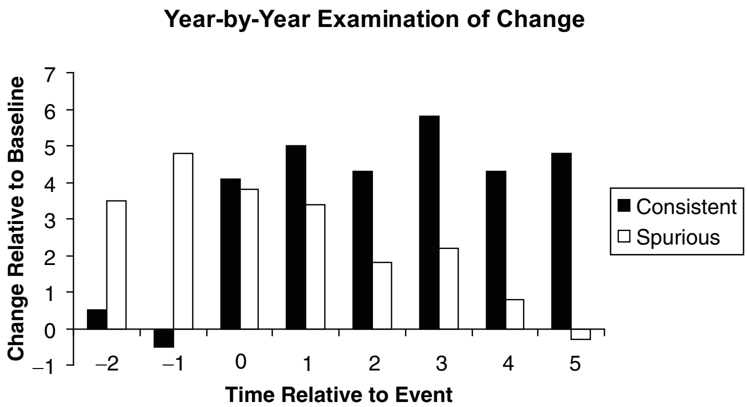

FIGURE 19.6. More detailed examination of timing of effects.

Figure 19.6 illustrates potential results from such an analysis. In one case, the findings are consistent with a causal interpretation because there is no systematic change in the outcome until the event occurs, after which there is a fairly consistent difference from the baseline. In the second case, the major change precedes the event, so we would likely have to judge that any association between the event and offending was spurious rather than causal.

Results from (19.12) may also be useful in guiding the specification of a more parsimonious model. For instance, a chart of the timing of changes such as Fig. 19.6 would be helpful for deciding whether the initial change in level is simultaneous or delayed (guiding whether to use X_{ti} or X_{t-1i} in the model), whether there is merely a change in level or in both level and slope (e.g., whether to use (19.8) or (19.11)), and whether any change in slope is likely linear or curvilinear (i.e., how to code $\text{Time}(X)_{ti}$ in (19.11)).

The utility of the model of (19.12) depends on the number and spacing of observations in the study. This approach would be most useful with many observations close in time, such as the event calendar approach developed by Horney and colleagues (see Chap. 15), which has often been used to gather 36 monthly observations for each respondent. In contrast, a panel study with only two or three time points could distinguish few time lags, and a study with observations spaced years apart is uninformative about the timing of any causal process operating on shorter time scales.

DIFFERENTIAL EFFECTS OF EVENTS

Criminologists also have offered many interesting hypotheses in which the effect of an event depends on some other variable that does not change over time. For instance, Moffitt (1993) argued that role transitions such as marriage and employment would have a greater impact for late onset offenders than for early onset offenders (a topic investigated in terms of propensity toward crime by Ousey and Wilcox 2007, and Wright et al. 2001). In this case, the impact of the event depends on an aspect of the individual's personal history. Differential effects could also arise in relation to other stable characteristics such as demographic variables or personality traits. A similar possibility is that the impact of an event might depend on some attribute of that event. That attribute could be some quality of the event, such as Sampson and Laub's (1993) argument that only high quality marriages and jobs will serve as turning points that reduce offending, or it could be the timing of the event, such as Uggen's (2000) finding that gaining employment brings crime reductions for older offenders but not younger offenders.

The general form of models that allow such differences in the effects of events is:

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + \beta_2 Z_i + \beta_3 X_{ti} \times Z_i + \beta_4 \text{Age}_{ti} + \beta_5 Z_i \times \text{Age}_{ti} + e_{ti} \qquad (19.13)$$

Z_i represents a time constant variable, which lacks the time-varying index t because it varies across people but not time. The parameters β_2 and β_5 allow for differences in the overall level and baseline age trend related to this variable, while β_3 captures variation in the effect of the event.[7]

Figures 19.7 and 19.8 illustrate the additional possibilities that arise in this model. Fig. 19.7 shows a hypothetical pattern of gender differences in the impact on crime of becoming a parent. As one would expect, the baseline age trends differ considerably between males and females, with males having a higher offense rate that declines more rapidly. If Z were coded as zero for females and one for males, then this pattern would correspond to a positive value for β_2 and a negative value for β_5. Because women typically bear the larger share of responsibilities for child-rearing, we might expect parenthood to have a greater impact on their lives, as shown here. With this coding for gender, the sizable decline in offending for females upon becoming a parent matches a large negative value for β_1, while the small decline in offending for males matches a positive value for β_3 (because they decline less than females).

Figure 19.8 shows a more complex example with differential change on both levels and slopes. In this case, the time constant variable is the age at which employment begins, and the figure indicates that years of full time employment brings less benefit to the earnings of people who begin full time employment younger rather than older. As we will see below, adjusting age trends for the occurrence or timing of the event has other useful implications.

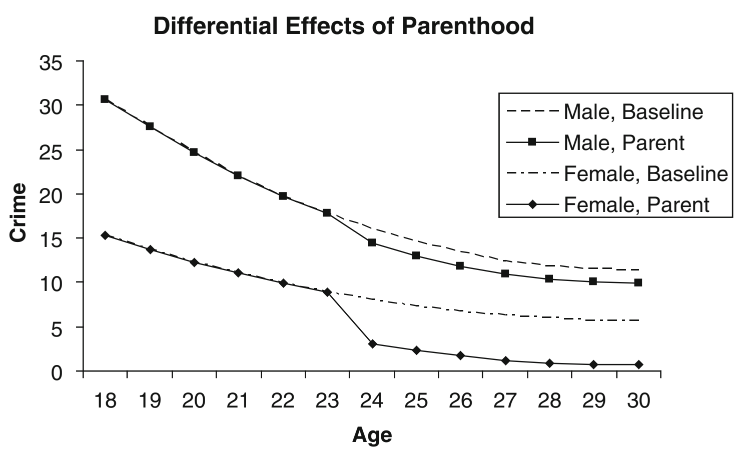

FIGURE 19.7. Differential effects of parenthood for males and females.

[7] To maintain a consistent focus on the relationships of events to within-individual change, it is necessary either to include the individual mean of the interaction term in the model as well, or else to subtract the individual mean from the event variable before forming its interaction with the time-constant variable.

FIGURE 19.8. Both effect and baseline varying by the age at which the event occurs: the effect of employment on crime.

Addressing the Possibility of Differing Age Trends

As discussed in an earlier section, models of the impact of events should allow for overall age trends in order to insure that results cannot be explained by the typical pattern of maturation. Analyses of differential effects based on (19.13) take this theme one step farther by allowing the age trends to differ in relation to the time stable variable of interest, such as the different baseline trends for males and females in Fig. 19.7. Doing so seems prudent in that it compares each individual to others who are more similar to themselves. Yet this also raises the broader question of what constitutes an appropriate baseline age trend for estimating the impact of an event.

It is useful to think of this issue in terms of the counter-factual model that dominates current thinking on causal inference (e.g., Winship and Morgan 1999). In this framework, the causal effect for each person is defined by the difference between the value of the outcome variable in the presence of the cause and its value in the cause's absence. The difficulty of causal inference is that we can never have the data needed to make this comparison. Because the cause cannot be simultaneously present and absent for the same person, at any given time we observe only one state of affairs or the other. Random assignment solves this problem by guaranteeing that, in the absence of a causal effect, the distributions of the outcome variable in the two conditions would differ only by chance.

For studying the effects of events, I have recommended that researchers routinely limit their analyses to within-person change and control for the overall age trend, each of which yields a considerable improvement over the basic regression model of (19.1). First, analyzing within-individual change narrows the comparison for observations in the presence of the cause to observations for the same person in its absence. Second, adjusting for the typical change between ages further refines the comparison to implicitly reflect a "difference in differences" comparing change for people who experience the event with change over the same ages for people who do not. Given the limitations of nonexperimental research designs, this is a relatively strong strategy because it rules out broad classes of plausible alternative explanations of findings. Specifically, the estimated effect cannot be due to any unchanging personal characteristics (whether measured or not) or to the typical pattern of maturation.

Of course, given an observational research design, even comparing people to themselves can never fully resolve the counterfactual dilemma. The data before and after the event are necessarily obtained at different ages and under different life circumstances, and we have no basis for arguing that at those occasions the outcome has identical distributions. As noted above, this leaves open the possibility of spurious effects due to omitted time-varying variables that could also contribute to differences in the outcome between the two time points, as well as the potential for reverse causality. These possibilities must be addressed by measuring and including the omitted variables most likely to be correlated with the event of interest and also to influence crime and by closely examining the timing of the effect of the event of interest (as in Fig. 19.6).

An additional challenge to causal inference when using this analytic approach is establishing an appropriate expectation for the change that would occur in the absence of the potential causal event, which is an essential component of its implicit counterfactual comparison. The models I have presented to this point largely rely on the overall pattern of maturation or change for this purpose, but we also need to address the possibility that people experiencing the event would not, on average, have followed this pattern. To date, the issue has received little attention, and it is an important direction for future methodological development. In the remainder of this section I offer some thoughts about the nature of this problem and potential strategies for addressing it.

The essence of the issue is the possibility that underlying developmental trends differ across individuals in such a way that the pattern of change for people who experience the event would be different from that of people who do not, even in the absence of the event. If not incorporated in the analysis somehow, this difference would create a bias in our estimate of the events' effect.

One possibility for addressing such differences in age trends is through interactions between age and individual characteristics. Equation (19.13) includes such an interaction to allow the underlying age trend to differ in terms of gender, and adding additional interactions with age would permit age trends to depend on a larger set of predictors. Suppose that these additional terms explained all of the reliable variation in age trends (evidenced by reduction in random variance components for the age trend), while there was little change in the estimated effect of the event. This pattern would enhance the plausibility of a causal effect by demonstrating that the result does not appear to be a spurious result of an undifferentiated specification of the baseline age trend that is inaccurate for this subpopulation.

A simpler and perhaps more effective means of establishing an appropriate baseline of expected change would be to focus more directly on who does and does not experience the event, and who does so earlier versus later. This strategy enhances the likelihood of taking into account the individual differences in patterns of developmental change that are most relevant to this event. Figure 19.8 illustrates this approach by making the baseline age trends (shown by the dashed lines) specific to people who experience the event at the same age.

A model to differentiate among people who experienced the event for longer and shorter periods of time would have some interesting and familiar features. The individual mean over time of the event variable, $\overline{X}_{\bullet i}$, indexes this variation, for instance, reflecting the proportion of time an individual was employed (to use the example of Fig. 19.8). Equation (19.14) allows age trends to depend on this variable:

$$Y_{ti} = \beta_0 + \beta_1 X_{ti} + \beta_2 \overline{X}_{\bullet i} + \beta_3 \text{Age}_{ti} + \beta_4 \overline{X}_{\bullet i} \times \text{Age}_{ti} + e_{ti} \tag{19.14}$$

Earlier in this chapter, a main effect for $\overline{X}_{\bullet i}$ appeared in (19.7), and in that case, it served to restrict the estimate of the event's effect to within-individual change by making that estimate independent from individual differences in average response level. Equation (19.14) adds the interaction of this variable with age as a way to make the estimated effect of the event independent from individual differences in the rate of age related change.[8]

To gain perspective on the causal model implicit in the data analysis approach presented in this chapter, it is useful to consider the multiple sources of information that contribute to the baseline of expected change. People who experience the event contribute to the overall baseline age trend through their waves of data before the event occurs. Individuals observed for at least two waves prior to the event provide information about the rate of age related change prior to the event (Bryk and Weisberg 1977). Using this information as the baseline of expected change entails projecting the trend beyond the available data. In the absence of other supporting information, such projections are always risky, and they are especially risky when projecting far into the future and when the trend is curvilinear, as age trends so often are.

People who do not experience the event provide a second source of information about expected age trends. Information from this group has the advantage that it extends throughout the period of study, including time during which others experience the event. Thus, it is free from the problem of projection beyond the data. The corresponding disadvantage of relying on this group as a baseline for change is that their age trend may differ from that of people who do experience the event.

The plausibility of viewing estimates from these models as causal effects will be considerably stronger if these two sources of information are in agreement about the expected pattern of change. Equation (19.14) provides a convenient vehicle for assessing whether they converge. For individuals who will experience the event, change over time before the event occurs is $\beta_3 \text{Age}_{ti} + \beta_4 \overline{X}_{\bullet i} \times \text{Age}_{ti}$, while for other individuals, change over time is $\beta_3 \text{Age}_{ti}$. Thus, β_4 captures the difference in age trend, and if its value is zero, then the baseline pattern of change does not depend on whether and for how long the event is experienced.[9]

What if β_4 is not close to zero, thus indicating differences in baseline age trends for people who do and do not experience the result? A causal interpretation of the effect of the event may still be plausible if allowing the baseline age trend to depend on the mean of the event variable does not alter the estimate of the effect (i.e., β_1 is little changed upon adding $\overline{X}_{\bullet i} \times \text{Age}_{ti}$ to the model). This pattern would suggest that, though baseline age trends are related to experiencing the event, they do not differ in a way that would account for the relationship between the event and the outcome variable.

The case for interpreting the estimated relationship as a causal impact of the event is weaker if adding this interaction alters the estimate. This pattern of findings would indicate that the estimated effect is sensitive to the set of individuals included in the analysis, which

---

[8] Halaby (2003) recommended using this interaction as a means for resolving a correlation between a time-varying covariate and a random effect for age, and I am suggesting that including it also enhances causal inference by adjusting for differing underlying age trends in relation to the event variable of interest. This is a close parallel to what is gained when we include the main effect for the individual mean on the event variable, as in (19.7). Doing so solves a technical statistical problem by removing any correlation between the time-varying covariate and the residual term for varying intercepts, while it serves the important methodological role of limiting the analysis to within-individual change.

[9] Data with curvilinear patterns of change would require the same comparisons for additional interactions of $\overline{X}_{\bullet i}$ with polynomial terms for age.

suggests that it would be wise to limit the analysis to the individuals for whom such comparisons are most justified. It would be plausible to do so through matching or weighting cases in terms of their propensity for experiencing the event (Rosenbaum and Rubin 1983), as several authors have recently pursued in studying relationships of events to crime (Haviland et al. 2007; King et al. 2007; Sampson et al. 2006).

In addition to these two sources of information, the baseline age trend is also affected by other constraints inherent in a specific statistical model. For instance, in most models I have discussed (including (19.14)), the event affects only the level of the outcome, and not its slope over time. This constraint implicitly assumes that the age-related change after the event follows the same pattern as that of people who either never experience the event or experienced it at a younger age. Accordingly, this constraint means that these postevent age trends also contribute to the baseline or expected age trend, while they would not (or at least would contribute less), if we removed the constraint by allowing the event to affect the slope (as in (19.11) and Figs. 19.4 and 19.8). Because such constraints bring more information to bear in estimating the effect of an event, they tend to increase statistical power and also provide a broader base for causal inference. At the same time, applying such constraints will distort the estimated effects if they are inconsistent with the data. Thus, it would be appropriate to test them by means such as assessing whether effects of events vary with age ((19.10)) or duration of the event ((19.11)).

CONCLUSION

In sum, the statistical models I have discussed provide flexible tools for using longitudinal data to examine the effects of events on crime. First, these models allow us to examine in detail the timing of these effects and to study differences in effects across groups of people and types of events. Furthermore, the models offer means of either testing for or eliminating many alternative explanations that compete with a causal interpretation of the estimated effect of the event. They allow us to address preexisting individual differences, overall developmental trends, the possibility that the supposed cause actually precedes the event, and the appropriateness of the baseline age trend as a basis for estimating the effect of the event.

One should not forget, however, that the basic longitudinal panel study has a passive observational research design, and it can never prove causality. Even so, this research design provides a great deal of information useful for eliminating or testing competing explanations, and thereby strengthening the plausibility of a causal interpretation of results. The techniques described in this chapter are meant to help researchers make the most of the information inherent in this research design. Successful application of this analytic approach has the potential to narrow the plausible alternative explanations to reverse causality and to other time-varying variables that are correlated with the one of interest. Controlling for time-varying variables and careful examination of the timing of the effect, as in Fig. 19.6, may provide evidence against those alternative interpretations as well.

There are, of course, other valuable approaches for studying causal effects, most notably, random assignment experiments, instrumental variable analyses (especially those using natural experiments or exogenous events as instruments), and propensity analyses (Winship and Morgan 1999). I suspect that life course criminologists will find the approach presented here especially useful, because they are most likely to find appropriate data on relevant topics in longitudinal panel studies (for examples, see Liberman 2007). Even so, confirmation of findings across methods would be especially valuable, and I encourage life course criminologists to be watchful for diverse opportunities to investigate effects of life events on crime.

REFERENCES

Agnew R (1992) Foundation for a general strain theory of crime and delinquency. Criminology 30:47–87
Agresti A (2007) An introduction to categorical data analysis, 2nd edn. Hoboken, NJ, Wiley
Akers RL (1977) Deviant behavior: a social learning perspective. Belmont, CA, Wadsworth
Allison PD (1990) Change scores as dependent variables in regression analysis. Sociol Methodol 20:93–114
Allison PD (2005) Fixed effects regression methods for longitudinal data using SAS. SAS Institute Inc, Cary, NC
Bryk A, Raudenbush SW (1987) Application of hierarchical linear models to assessing change. Psychol Bull 101:147–158
Bryk A, Weisberg H (1977) Use of the nonequivalent control group design when subjects are growing. Psychol Bull 84:950–962
Bushway S, Brame R, Paternoster R (1999) Assessing stability and change in criminal offending: a comparison of random effects, semiparametric, and fixed effects. J Quant Criminol 15:23–61
Campbell DT, Stanley JC (1966) Experimental and quasi-experimental designs for research. Chicago, Rand McNally
Curran PJ, Bollen KA (2001) The best of both worlds: combining autoregressive and latent curve models. In: Collins LM, Sayer AG (eds) New methods for the analysis of change. American Psychological Association, Washington, DC, pp 105–135
Esbensen F-A, Wayne Osgood D, Taylor TJ, Peterson D, Freng A (2001) How great is G.R.E.A.T.?: results from a longitudinal quasi-experimental design. Criminol Public Policy 1:87–115
Felson M (2002) Crime and everyday life: insights and implications for society, 3rd edn. Pine Forge Press, Thousand Oaks, CA
Greene WH (2000) Econometric analysis. Prentice Hall, Upper Saddle River, NJ
Halaby CN (2003) Panel models for the analysis of change and growth in life course studies. In: Mortimer JT, Shanahan MJ (eds) Handbook of the life course. Kluwer, New York, pp 503–527
Haviland A, Nagin DS, Rosenbaum PR (2007) Combining propensity score matching and group-based trajectory analysis in an observational study. Psychol Methods 12:247–267
Hay C, Forrest W (2006) The development of self-control: examining self-control theory's stability thesis. Criminology 44:739–774
Hirschi T (1969) Causes of delinquency. University of California Press, Berkeley, CA
Hirschi T, Gottfredson M (1983) Age and the explanation of crime. Am J Sociol 89:552–584
Hirschi T, Gottfredson M (1985) All wise after the fact learning theory, again: reply to baldwin. Am J Sociol 90: 1330–33
Horney J, Osgood DW, Marshall IH (1995) Criminal careers in the short-term: intra-individual variability in crime and its relation to local life circumstances. Am Sociol Rev 60:655–73
Jacobs JE, Lanza S, Osgood DW, Eccles JS, Wigfield A (2002) Changes in children's self-competence and values: gender and domain differences across grades one through twelve. Child Dev 73:509–527
Jang SJ (1999) Age-varying effects of family, school, and peers on delinquency: a multilevel modeling test of interactional theory. Criminology 37:643–685
Johnson DR (1995) Alternative methods for the quantitative analysis of panel data in family research: pooled time-series methods. J Marriage Fam 57:1065–1077
King RD, Massoglia M, MacMillan R (2007) The context of marriage and crime: gender, the propensity to marry, and offending in early adulthood. Criminology 45:33–65
Laub JH, Nagin DS, Sampson RJ (1998) Trajectories of change in criminal offending: good marriages and the desistance process. Am Sociol Rev 63:225–238
Laub JH, Sampson RJ (2003) Shared beginnings, divergent lives: delinquent boys to age 70. Harvard University Press, Cambridge, MA
LeBlanc M, Loeber R (1998) Developmental criminology updated. In: Crime and justice: a review of research, vol 23. pp 115–198
Liberman AM (ed) (2007) The long view of crime: a synthesis of longitudinal research. Springer, New York
Long JS (1997) Regression models for categorical and limited dependent variables. Sage, Thousand Oaks, CA

McClendon MJ (1995) Multiple regression and causal analysis.Peacock Publishers, Itasca, IL
Moffitt TE (1993) Adolescence-limited and life-course-persistent antisocial behavior: a developmental taxonomy. Psychol Rev 100:674–701
Nagin DS, Farrington DP, Moffitt TE (1995) Life-course trajectories of different types of offenders. Criminology 33:111–139
Nagin DS, Tremblay RE (1999) Trajectories of boys' physical aggression, opposition, and hyperactivity on the path to physically violent and non-violent juvenile delinquency. Child Dev 70:1181–1196
Nagin DS, Tremblay RE (2005) What has been learned from group-based trajectory modeling? examples from physical aggression and other problem behaviors. Ann Am Acad Pol Soc Sci 602:82–117
Osgood DW (2005) Sense of crime and the life course. Ann Am Acad Pol Soc Sci 602:196–211
Osgood DW, Rowe DC (1994) Bridging criminal careers, theory, and policy through latent variable models of individual offending. Criminology 32:517–554
Osgood DW, Smith G (1995) Applying hierarchical linear modeling to extended longitudinal evaluations: the boys town follow-up study. Eval Rev 19:3–38
Osgood DW, Wilson JK, Bachman JG, O'Malley PM, Johnston LD (1996) Routine activities and individual deviant behavior. Am Sociol Rev 61:635–655
Ousey GC, Wilcox P (2007) The interaction of antisocial propensity and life-course varying predictors of delinquent behavior: differences by method of estimation and implications for theory. Criminology 45:313–354
Paternoster R, Bushway S, Brame R, Apel R (2003) The effect of teenage employment on delinquency and problem behaviors. Soc Forces 82:297–335
Peterson Trond (1993) Recent advances in longitudinal methodology. Annu Rev Sociol 19:425–454
Piquero AR, Farrington DP, Blumstein A (2007) Key issues in criminal career research: new analyses of the cambridge study in delinquent development. Cambridge University Press, Cambridge, UK
Raudenbush SW (2001a) Comparing personal trajectories and drawing causal inferences from longitudinal data. Annu Rev Psychol 52:501–525
Raudenbush SW (2001b) Toward a coherent framework for comparing trajectories of individual change. In: Collins LM, Sayer AG (eds) New methods for the analysis of change,American Psychological Association, Washington, DC, pp 33–64
Raudenbush SW, Bryk AS (2002) Hierarchical linear models, 2nd edn. Sage, Newbury Park, CA
Rosenbaum PR, Rubin DB (1983) The central role of the propensity score in observational studies for causal effects. Biometrika 70:41–55
Rovine MJ, Molenaar PCM (2001) A structural equations modeling approach to the general linear mixed model. In: Collins LM, Sayer AG (eds) New methods for the analysis of change, American Psychological Association, Washington, DC, pp 65–96
Sampson RJ, Laub JH (1993) Crime in the making: pathways and turning points through life. Harvard University Press, Cambridge, MA
Sampson RJ, Laub JH, Wimer C (2006) Does marriage reduce crime? A counterfactual approach to within-individual causal effects. Criminology 44:465–508
Siennick SE, Osgood DW (2007) A review of research on the impact on crime of transitions to adult roles. In: Liberman AM (ed) The long view of crime: a synthesis of longitudinal research, Springer, New York, pp 161–190
Singer JD, Willett JB (2003) Applied longitudinal data analysis. Oxford University Press, New York
Slocum LA, Simpson SS, Smith DA (2005) Strained lives and crime: examining intra-individual variation in strain and offending in a sample of incarcerated women. Criminology 43:1067–1110
Snijders TAB, Bosker R (1999) Multilevel analysis: an introduction to basic and advanced multilevel modeling. Sage, Thousand Oaks, CA
Sutherland EH, Cressey DR (1955) Principles of criminology, 5th edn. J. B. Lippincott, Philadelphia
Thornberry TP (1987) Toward an interactional theory of delinquency. Criminology 25:863–891
Thornberry TP, Krohn MD (2005) Applying interactional theory to the explanation of continuity and change in antisocial behavior. In: Farrington DP (ed) Integrated developmental and life course theories of offending. Advances in criminological theory, vol 14. Transaction, Piscataway, NJ, pp 183–210
Uggen C (2000) Work as a turning point in the life course of criminals: a duration model of age, employment, and recidivism. Am Sociol Rev 67:529–546
Warr M (1993) Age, peers, and delinquency. Criminology 31:17–40
Winship C, Morgan SL (1999) The estimation of causal effects from observational data. Annu Rev Sociol 25:659–707
Wright BRE, Caspi A, Moffitt TE, Silva PA (2001) The effects of social ties on crime vary by criminal propensity: a life-course model of interdependence. Criminology 39:321–348