


**Center for Demography and Ecology
University of Wisconsin-Madison**

Event History Models for Life Course Analysis

Lawrence L. Wu

CDE Working Paper No. 2001-17



EVENT HISTORY MODELS FOR
LIFE COURSE ANALYSIS

Lawrence L. Wu
Department of Sociology and
Center for Demography and Ecology
University of Wisconsin-Madison

September 2001
[Revised February 2002]

Prepared for Jeylan Mortimer and Michael Shanahan (Editors), *Handbook of the Life Course*. Research funding from the National Institute of Child Health and Human Development (HD 29550) and research facilities provided under HD 05876 to the Center for Demography and Ecology, are gratefully acknowledged. Direct all correspondence to Lawrence L. Wu, Department of Sociology, 1180 Observatory Drive, University of Wisconsin-Madison, Madison, WI 53705, lwu@ssc.wisc.edu.

The questions posed by life course researchers often differ in fundamental ways from those posed by sociologists, developmental psychologists, or economists (Elder 1998; Mayer and Tuma 1990). For example, life course researchers often focus analytic attention on transitions marking adolescence or early adulthood, and the roles and statuses accompanying such transitions (Modell, Furstenberg, and Hershberg 1976; Hogan and Astone 1986; Shanahan 1999). Prototypical questions along these lines include whether certain social groups experience a more rapid transition to adulthood or whether the timing of such transitions (or the duration spent in selected life course statuses) has changed for successive cohorts (Winsborough 1980). Such an approach, as Mayer and Tuma (1990) note, implicitly conceives of social structure as arising out of individual experiences of varying duration, as opposed to more traditional views that see social structure in terms of collectivities of persons with particular fixed attributes (Blau 1977), as generated from relational networks (and resulting “structural holes”) among individual actors (White, Boorman, and Breiger 1976), or from the aggregate behavior of rational actors (Becker 1991).

As a result, life course researchers have stressed links between early and later life course events, roles, and attitudes (Elder 1999).¹ This focus, in turn, carries implications for the hypotheses typical of, and statistical models appropriate for, research on the life course, implying a wider range of questions and richer set of covariates—including how dynamic measures of past events, experiences, and other aspects of individual biography might shape current and future trajectories of behavior—than is typical of more traditional sociological research

A concern with the past, present, and future, including how biographical trajectories are shaped by larger social forces, has also led life course researchers to inquire more generally into the various temporal and dynamic processes shaping individual trajectories. As a result, life course researchers have come to reject, by and large, the notion that the life course can be understood simply as a process of unilinear aging (Mayer and Tuma 1990; Settersten and Mayer 1997).

¹This focus is not unique to the life course; indeed, precursors can be found in much work in the status attainment literature; see, e.g., Blau and Duncan (1967) and Sewell and Hauser (1975).

This has spurred life course researchers to be sensitive to the analytical importance of multiple dimensions of time—for example, age, duration in various statuses, and historical dimensions of time as measured by period or cohort.²

A final theme emphasizes interrelationships in the life course, two aspects of which are often stressed. One concerns the assertion, testable in principle, that domains such as work, marriage, childbearing, and emotional development, which are typically analyzed in isolation, in fact cannot be adequately understood without considering these domains as a unified whole. Another is the assertion, again testable in principle, concerning “linked” lives—that, for example, the events, behaviors, and outcomes experienced by one individual in a couple profoundly influences the course of events, behaviors, and outcomes experienced by a spouse.

In this chapter, I review methods relevant to life course research when event history data—that is, data on one or more discrete outcomes followed (at least in principle) in continuous time—are available. Several excellent monographs and textbooks (Allison 1985; Blossfeld, Hamerle, and Mayer 1989; Cox and Oakes 1984; Fleming and Harrington 1991; Hougaard 2000; Lancaster 1990; Tuma and Hannan 1984; Yamaguchi 1991) provide extensive coverage of relevant statistical models and issues. I review these issues, but also devote attention to how existing models speak to (or, in some cases, do not speak to) the types of questions most often posed by life course researchers.

Conceptualizing transitions between discrete life course statuses

Event history analysis is well-suited to an analysis of life course transitions.³ Indeed, the very concept of a transition is central both to the research on the life course and to the conceptual

²Another important and highly influential line of research has emphasized the long historical view, for example, the collective life trajectories of birth cohorts (Elder 1999; Mayer 1988), and the increasing differentiation of individuals from one another as a consequence of the elaboration of distinct institutional domains, on the one hand (Meyer 1986), and the increasing diversity of family and work experiences, on the other (Bumpass 1990; Spain and Bianchi 1996). For an overview of statistical issues that arise in this line of research, see Glenn (forthcoming).

³For an alternative view, see Abbott and Tsay (2000) and the resulting commentary (Levine 2000; Wu 2000; Abbott 2000.)

and statistical modeling of event histories. For concreteness, consider Figure 1, which sketches two ways a researcher might conceptualize a particular life course transition for women—that of the transition to single motherhood. Panel A of Figure 1 presents a simple and highly stylized conceptualization of this process in which women can occupy two statuses of interest—that of being a single mother and that of not being a single mother. All women thus begin life at birth in the status labelled “1” (not a single mother), with some women subsequently transiting to the status labelled “2.” Those who have become single mothers can exit this status as well, transiting back to the status labelled “1,” and so forth.

[Figure 1 about here]

Panel B presents an alternative conceptualization of this process that cross-classifies a woman’s marital status and number of children. Women begin life unmarried and with no children (status labelled “10”), and may then transit to subsequent statuses. (Child mortality is depicted in this panel by the leftward pointing arrows.) Note that there are now several possible transitions to single motherhood—the transition $10 \rightarrow 11$ consisting of a nonmarital first birth and the transitions $21 \rightarrow 11$, $22 \rightarrow 12$, and $23 \rightarrow 13$ representing changes to not married statuses of married women.

Because Panel B is an elaboration of Panel A, the corresponding models are nested within one another, with the model in Panel A equivalent to that in Panel B when a variety of behavioral assumptions are imposed on the model in Panel B. Equivalently, Panel A can be said to “pool” across the various transitions in Panel B, where, as in more standard contexts, such “pooling” assumptions can be tested empirically in ways that are formally equivalent to tests of pooling across race and ethnicity. Note that one could add further conceptual distinctions to the model in Panel B—one could, for example, distinguish between the childbearing of women who are never married, divorced, or widowed, or between the childbearing of women in cohabiting and marital unions compared to that of women in neither such union. Still, the point to be emphasized is the importance of considering alternative conceptualizations of the transitions of *theoretical* interest,

even when not all such transitions are observed or available in the data at hand.

Another important notion is the pool of individuals at *risk* of a particular transition. In Panel A, the risk set is straightforward—women at risk of a transition to single motherhood consist of those who are not single mothers, while those at risk of the other transition consist of single mothers—while in Panel B, different risk sets of women are distinguished. Under either conceptualization of the process, however, it is clearly important to restrict the sample analyzed to those at risk of a particular transition, with departures from this rule of thumb undertaken only when the researcher has clear substantive or theoretical grounds to do so.

Single transition for a homogeneous population

I now formalize ideas starting with the simple case of a single transition for a homogeneous population—for example, age-specific mortality in a population that is assumed to be behaviorally identical.⁴ An important issue is that the event in question (mortality, in this example) will often vary substantially with time (age in this example), with this pattern of time variation often exhibiting substantial nonlinearities. For example, age variation in mortality typically follows a so-called “bathtub” pattern, in which mortality is high at young and old ages, but low during the adolescent and adult years. Accounting for such patterns of time variation is of critical importance; in particular, estimation of other quantities of interest will, in general, be biased, sometimes substantially, if the underlying pattern of time variation is not accounted for adequately.

In modeling life course transitions, the quantity of fundamental interest is the so-called *hazard rate*, which can be defined variously as:

$$r(t) = \lim_{\Delta t \downarrow 0} \frac{\Pr(T \leq t + \Delta t | T \geq t)}{\Delta t} = \frac{f(t)}{1 - F(t)} = \frac{f(t)}{S(t)} \quad (1)$$

⁴Within an event history context, homogeneity does not imply that individuals in such a population experience the event of interest at the same time. As a rough analogy, consider a hypothetical population of coins in which the probability of heads is 1/2, where flipping 100 coins sampled from such a population will not yield 100 identical outcomes. A closer analogy is to radioactive material, in which individual atoms, even when chemically identical, will decay at different times. In the latter case, the *distribution* of event times can be shown to be exponential even the event time for any given atom cannot be predicted.

where t denotes time (for example, age), T denotes the random variable for the time of the event (age at death), and $f(t)$, $F(t)$, and $S(t) = 1 - F(t) = \Pr(T \geq t)$ denote the *density*, *cumulative distribution*, and so-called *survivor probability*, respectively.

The hazard rate provides the dependent variable of interest when modeling a single transition and is also a key building block for more complicated problems like that depicted in Panels A and B of Figure 1. It is, however, not an immediately intuitive quantity. Some insight can be gained by considering its component parts. For example, note that the units for $f(t)$ and $S(t)$ are percent per unit time and percent, respectively; thus, the unit for $r(t)$ is “per unit time” as in “the age-specific first birth rate per month.” Note also that because $f(t)$ and $S(t)$ are nonnegative quantities, $r(t)$ is also a nonnegative quantity that can assume any value between 0 and ∞ . Finally note that, as is the case in a logistic regression model or for the probability p governing a coin flip, the outcome—the probability p , the log odds $\log(p/1 - p)$, or the hazard rate, $r(t)$ —is not directly observed in the way a continuous outcome y_i is observed in a static linear regression model.

The quantity $\Pr(T \leq t + \Delta t | T \geq t)$ in (1) gives probability of having the event between time t and $t + \Delta t$, conditional on the event of interest not yet having occurred; hence, conditional on the population at risk, this quantity provides the probability that the event will occur between “now,” as indexed by t , and some time in the future, as indexed by $t + \Delta$. For events occurring in continuous time, it is desirable to define $\Pr(T \leq t + \Delta t | T \geq t)$ over all possible positive t ; this is done via the limit in (1), with the limit restricted to positive values of t so as to restrict intervals to those in the future. Combining these two parts—the limit and the conditional probability—yields the hazard rate, which is typically interpreted as the “risk” of an event, where risk refers to the “instantaneous” conditional probability that the event of interest occurs at time t .

To see that the three alternative definitions in (1) are equivalent, first consider the quantity $[\Pr(T < t + \Delta t | T \geq t)]/S(t)$, putting aside momentarily consideration of the limit and where for analytical convenience I assume an absolutely continuous $f(t)$. Recall from elementary probability theory that for $\Pr(B) > 0$ one can write

$$\Pr(A|B) = \frac{\Pr(A \text{ and } B)}{\Pr(B)}.$$

Let $A = T < t + \Delta t$ and $B = T \geq t$, in which case

$$A \text{ and } B = \{T < t + \Delta t\} \text{ and } \{T \geq t\} = t \leq T < t + \Delta t$$

Then

$$\begin{aligned} \Pr(T < t + \Delta t | T \geq t) &= \Pr(A|B) \\ &= \frac{\Pr(A \text{ and } B)}{\Pr(B)} \\ &= \frac{\Pr(t \leq T < t + \Delta t)}{\Pr(T \geq t)} \\ &= \frac{\Pr(t \leq T < t + \Delta t)}{S(t)}, \end{aligned} \tag{2}$$

recalling that $S(t) = 1 - F(t) = \Pr(T \geq t)$.

Note that the limit in (1) is absent from (2). To reintroduce the limit, recall from elementary probability theory that for an absolutely continuous density $f(t)$, one can write:

$$f(t) = \lim_{\Delta t \downarrow 0} \frac{\Pr(t \leq T < t + \Delta t)}{\Delta t}$$

Then applying the limit in (1) to (2) yields:

$$\begin{aligned} \lim_{\Delta t \downarrow 0} \frac{\Pr(T < t + \Delta t | T \geq t)}{\Delta t} &= \lim_{\Delta t \downarrow 0} \frac{\Pr(t \leq T < t + \Delta t)}{\Delta t S(t)} \\ &= \frac{f(t)}{S(t)}, \end{aligned}$$

which establishes the equivalence between the alternative definitions in (1).

A somewhat more intuitive quantity than $r(t)$ is the survivor probability $S(t)$; in population terms, this quantity for a single transition can be thought of as giving the proportion of the population that survives to time t without having experienced the event of interest. The hazard rate and survivor function are related according to:

$$S(t) = \exp\left[-\int_0^t r(s)ds\right] = \exp[-H(t)], \tag{3}$$

where the so-called *integrated hazard* is given by

$$H(t) = \int_0^t r(s)ds. \quad (4)$$

To gain further insight into the quantities $S(t)$, $H(t)$, $r(t)$, and $\log r(t)$, consider some hypothetical data on age at first marriage for three successive birth cohorts as presented in Figure 2. In this example, I constructed the curves for $r(t)$ from those for $\log r(t)$, with the three curves for $\log r(t)$ shifted vertically from one another by an additive constant:

$$\log r_A(t) = \log r_B(t) + c = \log r_C(t) + 2c.$$

The two upper panels present graphs of $r(t)$ and $\log r(t)$ for these three cohorts, with the solid curves in both panels lying uniformly above the dotted and dashed curves, and the dashed curve lying uniformly below the dotted curve. Exponentiating the above shows that the corresponding relationships for $r(t)$ take a multiplicative form, with

$$r_A(t) = r_B(t) \times \exp(c) = r_C(t) \times \exp(2c).$$

This is an example of the widely used *proportional hazard specification*, in which plots of $\log r(t)$ are parallel or in which the ratio of hazard rates is a constant that does not vary with t :

$$\frac{r_A(t)}{r_B(t)} = \exp(c), \quad \frac{r_A(t)}{r_C(t)} = \exp(2c), \quad \text{and} \quad \frac{r_B(t)}{r_C(t)} = \exp(c).$$

The quantities $\exp(c)$ and $\exp(2c)$ can be interpreted as *relative risks*. For example, in constructing Figure 2, I took $c = 0.5$; thus, for this hypothetical example, the age-specific “risk” of first marriage is $\exp(+0.5) = 1.65$ or 65 percent higher for the cohort A (solid curve), relative to that for cohort B (dotted curve), while the age-specific risk of first marriage for cohort C (dashed curve) is $\exp(-0.5) = .61$ or 39 percent lower relative to cohort B (dotted curve) and $\exp(-1.0) = .37$ or 63 percent lower relative to cohort A (solid curve).

[Figure 2 about here]

The bottom two panels of Figure 2 plot the corresponding survival probabilities and integrated hazard functions. Note in particular the inverse relationship between $r(t)$ and $S(t)$, with higher rates corresponding to lower survival probabilities. Similarly, because the solid curve for $r(t)$ lies uniformly above the dotted curve for $r(t)$, it then follows that the solid curve for $S(t)$ will lie uniformly below the dotted curve for $S(t)$. In the context of age at first marriage, then, cohort *A* experiences uniformly higher age-specific rates of entry into first marriage than cohorts *B* or *C* (upper panels for $r(t)$ and $\log r(t)$); likewise, the proportions remaining never-married (lower left-hand panel) are uniformly lower in cohort *A* than in cohorts *B* or *C*. Similarly, by age 40, 5.6, 17.4, and 34.6 percent of individuals in cohorts *A*, *B*, *C*, respectively, remain single and never-married. Thus, if cohorts *A*, *B*, and *C* represented successive birth cohorts of women, these results would indicate both a delay in age at first marriage and a greater propensity to forgo entry into first marriage for successive birth cohorts of women.

The four panels of Figure 2 illustrate how graphical plots of $r(t)$, $\log r(t)$, and $S(t)$ can convey useful information. For example, the two upper panels reveal age-graded differences in first marriage. The quantity $S(t)$ similarly provides information on how the proportion of the population remaining single and never-married—that is, surviving in the origin state—varies with age. Similar analyses can be conducted for successive birth cohorts, by race and ethnicity, or for cross-classifications by other characteristics.

As noted, Figure 2 provides an example of a proportional hazard specification. Although widely used by social scientists, it is important to emphasize that many of the questions posed by life course researchers in fact imply violations of this assumption. For example, one debate concerning the entry into first marriage is whether the behavior of successive cohorts of women is best understood as a “retreat” from marriage (see, e.g., Popenoe 1996; Gilder 1986) or whether these behaviors instead represent delayed but eventual entry from marriage (Oppenheimer 1997).

As this example shows, a proportional specification carries strong assumptions concerning this question—if it were invoked for successive cohorts of women, it would imply both delay and

retreat. Yet while retreat necessarily implies delay, the converse need not hold. Figure 3 presents a hypothetical example in which retreat—in the sense of increases in the proportions who never marry—need not follow from delayed marriage. In the two upper panels of Figure 3, the curves for $r(t)$ and $\log r(t)$ are shifted horizontally, with the peak age at entry into first marriage occurring at age 18, 20, and 22 for cohorts A , B , and C , respectively. The lower left-hand panel traces the consequences by age in the proportions who remain single and never-married. Differences in these proportions are large between 15 and 30, but narrow substantially at later ages. Thus under this hypothetical marriage regime, delay does not imply retreat.

[Figure 3 about here]

Nonparametric analyses of a single transition.

As Figures 2 and 3 suggest, much can be learned from exploratory analyses that permit visual inspection of the quantities $S(t)$, $H(t)$, $r(t)$, and $\log r(t)$. This is typically accomplished using *nonparametric estimators* of the survival probability (Kaplan and Meier 1958), integrated hazard (Aalen 1978; Nelson 1972), and hazard rate (Cox and Oakes 1984), where nonparametric refers to the lack of strong assumptions concerning the form of time dependence in the hazard rate.

As in the previous section, I begin by considering the case of a single transition in a homogeneous population. Let T_i denote the random variable for individual i 's time at the event of interest. Not all individuals may experience the event by the time of last interview, in which case these individuals are then said to be *right-censored*. As a result, representing data on the outcome requires a pair of variables, (t_i, δ_i) , where t_i is the realization of the random variable T_i and $\delta_i = 1$ if the event is observed for individual i and 0 if the outcome for i is right-censored.

The Kaplan-Meier (1958) estimator of the survival probability, which can be shown to be the nonparametric maximum-likelihood estimator for this quantity, is given by:

$$S_{\text{KM}}(t) = \prod_{\mathcal{R}(s): s \leq t} \left[1 - \frac{d(s)}{\#\mathcal{R}(s)} \right], \quad (5)$$

where the product is taken over all individuals at risk of the event at time t , $d(t)$ denotes the number of individuals with events at time t , $\mathcal{R}(t)$ is the set of individuals at risk of the event at time t , and $\#\mathcal{R}(t)$ denotes the number of individuals at risk at time t . The Nelson-Aalen estimator (Nelson 1972; Aalen 1978), which can also be motivated as the nonparametric maximum-likelihood estimator for the integrated hazard, is given by:

$$H_{\text{NA}}(t) = \sum_{\mathcal{R}(s):s \leq t} \frac{d(s)}{\#\mathcal{R}(s)}. \quad (6)$$

For technical statistical reasons, no sensible nonparametric maximum-likelihood estimator of the hazard rate exists, but one can nevertheless form nonparametric estimators with good properties using the classic demographic life table or methods for nonparametric density estimation. For additional details, see, for example, Allison (1995), Cox and Oakes (1984), or Wu (1989).

Figure 4 presents nonparametric estimates of the two most intuitive quantities, $S(t)$ and $r(t)$, for age at first marriage for white and black women in the 1980 Current Population Survey (CPS). The upper panel, which presents Kaplan-Meier estimates of $S(t)$, shows that $S_{\text{KM}}(t)$ yields a monotonically declining step function, with values varying between 1 and 0. For these cohorts of women, the curves decline steeply between ages 20 and 40; roughly 64 percent of white women were never married by age 20, with about 7 percent never married by age 40. By contrast, about 69 and 17 percent of black women were never married by ages 20 and 40, respectively. The middle panel of Figure 4, which presents smoothed nonparametric estimates of $r(t)$, shows that black/white differences in the transition to first marriage are indeed most pronounced between ages 20 and 40. The lower panel, which presents smoothed nonparametric estimates of $r(t)$, shows that black/white differences are not constant with age, but rather widen during early and middle adulthood, and then narrow at later ages. This panel makes clear that proportionality appears to be violated for white and black patterns in age at first marriage; note also that black/white first marriage rates do not appear to follow a pattern of “pure” delay such as that exhibited in Figure 3.

These examples also illustrate the prototypical steps in how one might conduct exploratory

analyses of event history data. In particular, observed characteristics of individuals can be used in conjunction with the nonparametric methods outlined above—for example, by classifying individuals using observed values of a discrete covariate or by discretizing the values of a metric covariate, followed by visual inspection of graphical displays such as those in Figure 4, with such analyses providing, roughly speaking, an event history equivalent to exploratory methods using cross-tabulations or comparisons of group means.

[Figure 4 about here]

Parametric models for a single transition.

To this point, our discussion has focused on a single transition for a homogeneous population; thus, we have proceeded by and large without consideration of right-hand-side covariates. In this section, we consider how one might obtain estimates of the association of observed covariates with the hazard rate $r(t)$. To simplify matters, we continue to focus on a single transition.

The most common parametric specification for covariates is the so-called *proportional hazard* model given by:

$$r[t|\mathbf{x}_i(t)] = q(t) \exp[\mathbf{x}_i(t)] \quad (7)$$

or

$$\log r[t|\mathbf{x}_i(t)] = \log q(t) + \mathbf{x}_i(t), \quad (8)$$

where i indexes individuals and $\mathbf{x}(t)$ denotes a vector of observed covariates, some of which may be time-varying. A key assumption in (7) and (8) is that time-variation in $\log r(t)$ is captured by a single function—the so-called baseline hazard $\log q(t)$ —with all additional heterogeneity across individuals captured by the additive effects of the covariates $\mathbf{x}_i(t)$.

The types of covariates allowed by (7) and (8) span a remarkably wide range; indeed, one can include any aspect of an individual's history up to time t (Aalen 1978; Tuma and Hannan 1984). Thus, event history models let life course researchers examine an unusually rich set of covariates in ways that capture changes in the social and economic contexts of individuals that

might influence the outcome of interest. One can, for example, investigate the effect of earlier life events, current social circumstances, cumulative experiences, and exposure to particular statuses, which allows researchers to contrast effects of past, current, and cumulative experience on current behavior (Wu and Martinson 1993; Wu and Thomson 2001). See also Wu (1996), who compared three alternative effects of income in an individual's family of origin—a simple measure of income at time t , permanent and transitory measures of income, and measures of income level and income change. One can also incorporate future predictions (e.g., projected income, marriage market opportunities, or characteristics of potential mates; see, e.g., Dechter 1992) when the predicted values of such covariates are obtained from models using current and past covariate information for a given individual. As Mayer and Tuma (1990) note, this flexibility meshes well with a central presupposition of many life course researchers, which is that an individual's social context can vary considerably over time in ways rarely reflective of some stable social equilibrium. If so, then a central analytical task is to capture central features of this over-time variation for individuals.

Several measurement issues arise when incorporating time-varying covariates into (7) and (8). For example, if a time-varying covariate $x_i(t)$ is discrete and gathered as an event history, the analyst can determine the value of $x_i(t)$ at all observed times t during which an person i is at risk. In other circumstances, the value of $x_i(t)$ may be difficult to determine at all possible times. Examples include covariates such as income, work hours, expenditures, depression scales, or attitudes that can vary from one moment to another but which will often be measured only sporadically. In such cases, a commonly invoked and analytically convenient assumption is that the value of such a covariate is constant between measurements.⁵

The popularity of proportional hazard models stems in part from practical and theoretical considerations. The asymptotic properties of these models are well understood under a variety of conditions (see, e.g., Cox and Oakes 1984; Fleming and Harrington 1991), including quite general

⁵For alternative specifications relaxing this constancy assumption, see, e.g., Tuma and Hannan 1984. Note also that the plausibility of such an assumption will depend not only on the number and frequency of measurements, but also on the temporal variation in the covariate relative to the frequency of measurements.

conditions on the distribution of censoring times. Furthermore, under proportionality, covariates have linear effects on the logarithm of $r(t)$; hence, intuitions from ordinary and logistic regression carry over in a straightforward way to the proportional hazard model. Empirically, proportionality is often adequate in that estimated coefficients under proportionality are similar to those when proportionality is relaxed; this can often hold when the observation period is short relative to the mechanisms that generate variation over time in the effect of a covariate.

Proportionality is nevertheless a strong assumption and violations can occur empirically. As noted above, Figure 4 provides an empirical example where proportionality appears suspect; indeed, visual displays like those in the lower panel of Figure 4 provide a useful exploratory way to check proportionality. Fortunately, the models in (7) and (8) can be generalized to incorporate nonproportional effects of covariates; indeed, by appropriate definition of time-varying covariates, one can adopt standard software to estimate certain nonproportional models. Recall that when proportionality is violated, the effect of a covariate on the logarithm of the hazard rate will not be an additive constant, but rather will vary with t . Thus, one way to relax the proportionality assumption is to code a set of time-varying dummy variables that represent an exhaustive and mutually exclusive partition of the observation period. Interacting these dummy variables with a covariate x hypothesized to have nonproportional effect then yields a specification that models the effect of x as a step function of time. Other approaches suggested in the literature include a log multiplicative specification (Xie 1994), piecewise linear splines (Wu and Martinson 1993), and a local likelihood approach that makes few assumptions about the form of time-variation in covariate effects (Wu and Tuma 1990).

In practice, assumptions concerning the baseline hazard $q(t)$ are typically of greater concern than possible nonproportionalities in the effects of covariates; very roughly speaking, this often occurs when the baseline hazard accounts for a greater proportion of the observed variation in $r(t)$ than do the covariates $\mathbf{x}(t)$, as is often true in practice. A popular choice is a model due to Cox (1972), in which the baseline hazard $q(t)$ is allowed to be an unknown and unspecified function of

time. Under this model, maximum likelihood estimation is not possible, but the method of partial likelihood (Cox 1975) can be shown to yield consistent and asymptotically efficient estimates of the parameters \mathbf{b} under quite general conditions, including mild assumptions concerning the distribution of censoring (Andersen and Gill 1982).

While the Cox model yields estimates of the parameters \mathbf{b} for the effects of covariates, it does not provide any direct estimate of the baseline hazard function $q(t)$. Knowledge of $q(t)$ is often unnecessary for a number of important analytical purposes; for example, comparisons of control and treatment groups in medical trials evaluating the efficacy of a new drug or treatment in which the outcome is mortality from a specific form of cancer. However, many of the questions routinely posed by life course researchers require knowledge of $q(t)$; examples include comparisons across appropriately defined cohorts of individuals in the pace of childbearing, the timing of entry into marriage, or median time to divorce. Although estimates of $q(t)$ can be recovered in the Cox model, obtaining inferences about $q(t)$ under the Cox model is more difficult. As a result, it can be useful to consider parametric alternatives to the Cox model.

Various parameterizations for $q(t)$ have been proposed, including the exponential, Weibull, Gompertz, Makeham, log logistic, log Gaussian, Hernes, sickle, and Coale-McNeil models (Tuma and Hannan 1984; Blossfeld, Hamerle, and Mayer 1989; Wu 1990). Sometimes theory provides grounds to motivate a particular choice, but more often practical considerations (e.g., software availability) underlie these choices. Unfortunately, estimated effects of covariates can sometimes vary markedly across different functional forms, complicating matters for the analyst.

One reason for this sensitivity is that the models vary markedly in their specification of time variation in the baseline hazard. For example, some models assume that the baseline hazard increases or decreases monotonically (e.g., the Weibull, Gompertz, and Makeham models), while other models assume a unimodal shape for the baseline hazard, with the rate rising and then declining (e.g., the log logistic, log Gaussian, Coale-McNeil, Hernes, and sickle models). In addition, some models yield a distribution of event times that integrates to unity, implying that

all individuals will experience the event of interest if observed for a sufficiently long period of time (e.g., the exponential, Weibull, Makeham, log logistic, and log Gaussian models), while other models can, in some cases, yield a so-called defective distribution in which some individuals will never experience the event of interest, even if observed for an arbitrarily long time. Note that a defective distribution of event times is often substantively plausible; examples include marriage or sexual initiation in which some individuals, for example, those who have taken vows of celibacy or chastity, may never marry; parity-specific fertility, in which some individuals may never proceed to, say, a fifth birth; and residential moves, where some individuals may live all their lives in the residence in which they were born. Similarly, there are other instances in which a nondefective distribution of event times is desirable, with a classic example being human mortality.

As a result, many researchers use models that mimic the Cox model in the sense of providing a flexible functional form for the baseline hazard. One popular and easily implemented parametric alternative to the Cox model models the baseline hazard as a piecewise constant function—that is, as a step function of t . More formally, consider P time intervals, $(0, \tau_1]$, $(\tau_1, \tau_2]$, \dots , $(\tau_{P-2}, \tau_{P-1}]$, $(\tau_{P-1}, \infty]$, where the τ_p are prespecified by the analyst; then let $q(t)$ be defined by a series of constants on these intervals, i.e.,:

$$q(t) = \begin{cases} \exp(\lambda_1), & t \in (0, \tau_1], \\ \exp(\lambda_2), & t \in (\tau_1, \tau_2], \\ \dots & \\ \exp(\lambda_P), & t \in (\tau_{P-1}, \infty]. \end{cases} \quad (9)$$

or, equivalently,

$$\log q(t) = \begin{cases} \lambda_1, & t \in (0, \tau_1], \\ \lambda_2, & t \in (\tau_1, \tau_2], \\ \dots & \\ \lambda_P, & t \in (\tau_{P-1}, \infty]. \end{cases} \quad (10)$$

The resulting piecewise exponential model for the baseline hazard is easily implemented, for example, by defining P time-varying dummy variables corresponding to the P time intervals.⁶

⁶Note that (7) and (8) lack a time-invariant constant term—what would be the intercept in a linear regression. One can retain such a constant term by omitting one of the P intervals from estimation, in which case the estimates λ_p and $\exp(\lambda_p)$ provide contrasts with respect to the overall constant term.

A slight variant of the above lets $\log r(t)$ vary linearly within intervals, with the linear segments splined to yield a continuous function:

$$\log q(t) = \begin{cases} \lambda_1 + \gamma_1 t, & t \in (0, \tau_1], \\ \lambda_2 + \gamma_2 t, & t \in (\tau_1, \tau_2], \\ \dots & \\ \lambda_P + \gamma_P t, & t \in (\tau_{P-1}, \infty]. \end{cases} \quad (11)$$

subject to the $P - 1$ equality constraints

$$\begin{aligned} \lambda_1 + \gamma_1 \tau_1 &= \lambda_2 + \gamma_2 \tau_1 \\ \lambda_2 + \gamma_2 \tau_2 &= \lambda_3 + \gamma_3 \tau_2 \\ &\dots \\ \lambda_{P-1} + \gamma_{P-1} \tau_{P-1} &= \lambda_P + \gamma_P \tau_{P-1}, \end{aligned} \quad (12)$$

yielding the so-called splined piecewise Gompertz model for the baseline hazard.

Both the piecewise constant function in (11) and the piecewise linear spline in (12) provide very flexible specifications for $\log q(t)$; note, for example, that both can accommodate a variety of shapes for the baseline hazard, including monotonically increasing, monotonically decreasing, unimodal, or multimodal patterns of time variation. Note that, given P prespecified time intervals, the piecewise constant specification in (11) uses P degrees of freedom, while the piecewise linear spline in (12) uses $P + 1$ degrees of freedom.⁷

Table 1 presents estimates from three proportional hazard models using data on white women from the 1980 CPS. We contrast estimates from a Cox model, a piecewise constant, and piecewise linear spline model; for the latter two models, we have specified three time intervals, corresponding to ages 12 and 18, 18 and 25, and 25 and older. Note that these data provide large samples but relatively few covariates; hence, Table 1 reports estimates only for respondent's years of schooling completed by time of interview (discretized into 0–11, 13–15, and 16 or higher, with 12 years the omitted category) and year of birth (1930–49 and 1950 or later, with 1929 or earlier the omitted category). Estimates agree closely across the three models, with estimates from the Cox

⁷The piecewise linear spline uses $2P$ parameters but then invokes $P - 1$ equality constraints, thus requiring $2P - (P - 1) = P + 1$ free parameters.

and three-period linear spline model agreeing particularly closely. For the Cox and three-period linear spline models, white women with 11 or fewer years of schooling completed have a 28 percent ($\exp(.25) = 1.28$) higher rate of first marriage than white women with 12 years of schooling, while white women with some college education have a 30 percent ($\exp(-.35) = .70$) lower rate of first marriage than white women with 12 years of schooling.

[Table 1 about here]

Results from these models also show that the piecewise linear spline specification for $q(t)$ provides a substantially better fit to the data than the piecewise constant specification. Table 2 compares log likelihood values corresponding to four models for these data using a piecewise constant baseline with and without covariates, and a piecewise linear spline baseline with and without covariates. The lower panel of Table 2 provides two sets of χ^2 comparisons that provide tests of adding the five covariates in Table 1 and for modeling $\log r(t)$ using three constants vs. three linear splines. Adding covariates yields a χ^2 increment in fit of around 4,300, while allowing $\log r(t)$ to vary linearly yields a χ^2 increment in fit around 13,500. This is a typical result when analyzing life course data, with careful modeling of the baseline hazard often yielding much more substantial improvements in fit than the introduction of covariates. Intuitively, such large increments in fit are often observed because typical life course outcomes exhibits substantial within-individual time variation, with this time variation often more substantial than the variation observed across individuals. As a consequence, it is generally advisable to devote careful modeling attention to the form of time-variation in $q(t)$ to insure that conclusions about other parameters are not biased by incorrect or inappropriate assumptions about time dependence in $q(t)$.

[Table 2 about here]

The proportional models in equations (7)–(12) can be easily extended to accommodate multiple dimensions of time. Consider, for example, the transition between a first and second child as depicted in Panel B of Figure 1. Empirically, the second birth rate varies less with the age of

a woman than with the duration since first birth, yet age variation in the second birth typically cannot be ignored. Because of examples like this, analysts often have sound reasons to extend (7)–(12) to multiple dimensions of time. Let t denote age and u the duration since a first birth; then a straightforward extension of the proportional hazard model in (7) is:

$$r[t, u | \mathbf{x}_i(t)] = q_1(t)q_2(u) \exp[\mathbf{x}_i(t)], \quad (13)$$

or, equivalently,

$$\log r[t, u | \mathbf{x}_i(t)] = \log q_1(t) + \log q_2(u) + \mathbf{x}_i(t). \quad (14)$$

Note that (13) and (14) yield an age- and duration-specific model of second births under the assumption that the second birth rate is separable into two components, $q_1(t)$ and $q_2(u)$ (Lillard 1993; Wu and Martinson 1993). As for a single time dimension, one can use numerous parameterizations for $q_1(t)$ and $q_2(u)$, including those in (10) or (11).

It has long been recognized that linear regression models can be unidentified when controlling for multiple dimensions of time, with the classic example being simultaneous linear terms for age, period, and cohort (Glenn forthcoming). Although t and u will covary in strong ways for each individual in the sample, identification in (13) and (14) is often possible because $q_1(t)$ and $q_2(u)$ are typically highly nonlinear functions of t and u , with these nonlinearities helping to identify model parameters. For example, Wu and Martinson (1993) present models that control duration, age, period, and cohort; however, identification will in general become increasingly problematic as the number of time dimensions or parameters used to model the baseline functions increases. Given these issues, one sensible procedure is to identify, theoretically or empirically, those time dimensions which induce the greatest variation in the underlying hazard rate. One can then invest greater modeling effort for the “primary” time dimensions, for example, using the specification in (11), and less effort for “secondary” time dimensions, for example, by using the specifications in (10) or (11) coupled with relatively widely-spaced intervals.

A wide class of diagnostics for the above models can be obtained using so-called martingale residuals (see, e.g., Fleming and Harrington 1991, Chapter 4). Graphical displays of such residuals

can be used to assess the influence of particular observations and to check assumptions concerning proportionality and the functional form of covariate effects (i.e., are effects linear in x or in $\log x$) in ways analogous to diagnostic residual displays in linear regression.

Multiple origin and destination states

Thus far, we have formalized issues for a single transition, but the models discussed above generalize in straightforward ways to more complicated processes such as those depicted in Panel B of Figure 1. To simplify details, let us return to single homogeneous population but generalize the above to multiple origin and destination states. Let j and k index the origin and destination states, respectively; then let the *transition rate* $r_{jk}(t)$ be defined as

$$r_{jk}(t) = \lim_{\Delta t \downarrow 0} \frac{\Pr(T_{jk} \leq t + \Delta t | T_{jk} \geq t)}{\Delta t} = \frac{f_{jk}(t)}{1 - F_{jk}(t)} = \frac{f_{jk}(t)}{S_{jk}(t)}. \quad (15)$$

Thus, (15) generalizes (1) by representing each transition $j \rightarrow k$ by a unique transition rate $r_{jk}(t)$.

The generalization of the survivor probability to multiple origins and destinations involves some subtle but important shifts. Returning momentarily to the case of a single event, recall that the survivor probability has two equivalent interpretations: (a) as the probability of not yet having experienced the event of interest, and (b) as the probability of remaining in the origin state. When multiple destination states exist, (a) and (b) will, in general, differ; in addition, the interpretation of (a) is complicated by an identifiability issue when so-called competing risks are present.

To make issues concrete, consider Panel A of Figure 1, in which there is only one transition out of each origin status. In this case, matters reduce to the case for a single transition, conditional on status at origin. By contrast, in Panel B of Figure 1, each origin state has multiple destination states—for example, from the origin state 01, there are two possible transitions, $01 \rightarrow 11$ and $01 \rightarrow 20$; while from the origin state 21, there are three possible transitions, $21 \rightarrow 11$, $21 \rightarrow 20$, and $21 \rightarrow 22$. When individuals in an origin state are subject to multiple destination states, they are said to be subject to *competing risks*. When competing risks are present, it can be shown that

the interpretation given in (a) is unaffected but that in (b) must be modified in ways detailed below.

To formalize these ideas, suppose that individuals observed in origin state j are subject to multiple destination states, indexed by $k = 1, \dots, K_j$, where $K_j > 1$. Let $r_{jk}(t)$ denote the K_j transition rates corresponding to each of these transitions; then the probability of surviving to time t in state j , is given by

$$S_j(t) = \exp\left[-\int_0^t \sum_{k=1}^{K_j} r_{jk}(s)ds\right] = \exp\left[-\sum_{k=1}^{K_j} \int_0^t r_{jk}(s)ds\right] = \prod_{k=1}^{K_j} \exp\left[-\int_0^t r_{jk}(s)ds\right]. \quad (16)$$

Thus, the second interpretation of $S_j(t)$ —the probability of surviving in origin state j —is identical to that for a single transition, except that $S_j(t)$ conditions on origin state j —that is, it refers to the survival to time t of those individuals who have not exited the origin state j by time t .

Now consider the probability in (a)—that of not having experienced the event of interest. When competing risks are present, there is not a single event, but rather multiple events, that must be considered when accounting for exits from an origin state j . Consider a classic example in which age-specific mortality in humans is classified by cause of death—for example, deaths due to (1) cardiovascular disease, (2) cancer, (3) homicides and other acts of violence, (4) accidents, and (5) a residual category for all other causes of death. Distinguishing between multiple types of events gives rise to additional complications; in particular, the interpretation of $S_{jk}(t)$ and $r_{jk}(t)$ will in general differ from more familiar quantities such as the proportion in the population observed to experience the event k conditional on origin state j .

To see this, consider two birth cohorts followed until death and suppose that cohort members are identical in all respects except that deaths due to cardiovascular disease have been eliminated in cohort A but not in cohort B . It then follows that, in cohort A , if one cause of death is eliminated, the proportion in A experiencing other causes of death will necessarily increase. Let p_k^A and p_k^B denote the proportion of deaths of type k that occur in cohorts A and B ; then if cohorts A and B are identical save for $p_{A1} = 0$, it will nevertheless follow that $p_k^A \geq p_k^B$ for all remaining causes of death. The researcher only possessing estimates of p_k^A and p_k^B will observe

that $p_{Ak} \geq p_{Bk}$ for $k \neq 1$, and hence might be tempted to conclude that mortality in A and B differ fundamentally, with mortality from causes other than cardiovascular disease systematically higher in cohort A than in B . Yet by construction, the two cohorts have identical mortality risks save for the elimination of deaths from cardiovascular disease in cohort A . This apparent paradox would be avoided were comparisons based on the quantities $r_k(t)$ or $S_k(t)$. Thus, sufficiently large samples would reveal that $r_1^A(t) = 0$ while $r_1^B(t) > 0$ —mortality from cardiovascular disease is eliminated in cohort A but not B —but that mortality from other causes are otherwise identical, i.e., that $r_k^A(t) = r_k^B(t)$ for $k \neq 1$.

Turning this example on its head makes clear that, under competing risks, the $S_{jk}(t)$ cannot be interpreted as if they provided the proportions of those in origin state j who experience the event j , a statement that holds even when censoring is absent. Rather, the $S_{jk}(t)$ should be interpreted as giving the proportion in an origin state who would have experienced the j th transition *were all other competing transitions to be eliminated*. Note, moreover, that the plausibility of this interpretation rests heavily on this independence assumption.⁸ Violations of this assumption would include situations in which, say, those who are observed to die of one chronic condition—for example, cardiovascular disease—differ systematically in other ways that lead them to have higher (or lower) risk of another chronic condition. A difficulty is that the independence assumption under competing risks has been shown to be nonidentifiable in the sense that one cannot obtain formal tests of this assumption (Tsiatis 1975).⁹

⁸When covariates are included in the model, this assumption becomes one of conditional independence, i.e., the competing risks are assumed to be independent conditional on the covariates $\mathbf{x}(t)$. This assumption is similar to the so-called “irrelevance of alternatives” in a multinomial logistic regression. It is possible to state conditions that are slightly weaker than full independence for the competing risk model; however, such conditions carry little practical import. See Cox and Oakes (1984) for details.

⁹Consider the promotion of assistant professors in an academic department. If some junior faculty depart in anticipation of nonpromotion, then simply distinguishing between two sorts of events (promotion versus departure) will not correct the upward bias in estimates of tenure rates. Allison (1995) suggests a simple procedure to assess the sensitivity of estimates to such a possibility. Consider two situations, one in which all departing junior faculty would have in fact been fired at a time ϵ after they are observed to have departed, and another in which all departing junior faculty would have been promoted had they remained. Estimates under these two behavioral extremes can be used to construct Manski-type bounds for the usual naive estimate (Manski 1995).

Parametric models for multiple transitions.

Construction of parametric models for the $r_{jk}(t)$ in (15) is straightforward, with the underlying issues similar to those for modeling a single transition. For example, a proportional hazard specification incorporating covariates will be given by:

$$r_{jk}[t|\mathbf{x}_{jk}(t)] = q_{jk}(t) \exp[\mathbf{b}_{jk}\mathbf{x}_{jk}(t)], \quad (17)$$

where $q_{jk}(t)$ denotes the baseline for the transition from state j to k and $\mathbf{x}_{jk}(t)$ denotes a vector of (possibly time-varying) covariates. The jk subscript on \mathbf{x} emphasizes that one can specify different covariates across transitions and that the effects themselves will differ across transitions. In practice, however, researchers often employ the same set of covariates across the multiple transitions to aid the substantive interpretation of the coefficients \mathbf{b}_{jk} .

One can estimate (17) using a suitable generalization of the Cox model to multiple origin and destination states. Under this specification, the $q_{jk}(t)$ are assumed to be arbitrary unspecified functions of t that vary in arbitrary but unspecified ways across transitions. If the Cox model is not used, then (17) will require parameterizing the $q_{jk}(t)$. Guidance for these parameterizations can be obtained from suitable generalizations to the nonparametric methods discussed above. The resulting patterns can then be used to select a particular parameterization of the $q_{jk}(t)$, with the underlying issues essentially identical to those for a single transition. Multiple dimensions of time can also be handled in ways similar to those discussed for a single transition.

Estimation of (17) will, in general, yield a very large number of parameter estimates—there will, in general, be coefficients for each covariate and transition pair. For substantive and interpretive parsimony, one may wish to determine if the effect of a covariate is similar across selected transitions—for example, for the transitions depicted in Panel B of Figure 1, if income effects are similar for third and higher-order marital births. Such hypotheses can be evaluated using log likelihood ratio or BIC tests under equality constraints on the appropriate parameters across transitions. Suitable extensions of this same idea can be used to determine if one can “pool”

across transitions—for example, if, in Figure 1, model fit is not substantially degraded under the more parsimonious model in Panel A relative to the more complex model in Panel B.

Unobserved heterogeneity.

As noted at the outset of this chapter, life course theorists often assert that one must approach the life course holistically—that domains such as work, marriage, childbearing, and emotional development, typically analyzed in isolation, cannot be understood adequately without considering these domains as one unified whole. The models considered to this point help sharpen this idea. Take, for example, Panel B of Figure 1, which depicts transitions in two domains—marriage and childbearing. One can show that the application of the nonparametric estimator in (5) to these transitions will reproduce the distribution of individuals, both across the multiple statuses and over time. Extensions of these methods to incorporate covariates—the Cox model or the models in (9)–(12)—will likewise also reproduce the across-state and across-time distribution of individuals when proportionality holds. This implies, then, that the importance of this assertion lies not at the level of aggregate distributions, but rather at the level at which individual behaviors are modeled.

One approach formalizing this idea lets events in one life course domain affect transition rates in another domain. The models considered thus far accommodate this type of dependence by letting the researcher condition on any aspect of an individual's past history—including the individual's trajectory or past history in another life domain—when modeling $r_{ijk}(t)$. Thus, measures constructed from a person's trajectory of work and labor force experience up to time t can be used as right-hand-side covariates in modeling risks at time t of transitions in the realms of fertility, marriage, divorce, retirement, and so forth. Another possibility, often raised by economists, is that an individual's decisions about work, fertility, and marriage reflect an attempt to maximize utility across these joint spheres. If all relevant decision inputs are observed, then one can condition on suitable measures of these inputs, in which case no new issues arise. However, when some relevant inputs are unobserved, a number of subtle issues arise in ways that are

substantially more troublesome in a hazard regression context than in a static linear regression.¹⁰

Recall that a standard result for linear regression is that if an unobservable is uncorrelated with a covariate x of interest, then the OLS estimate of the effect of x will be unbiased although not optimally efficient. No analogous result holds in the hazard regression context; in particular, even if an unobservable z is initially uncorrelated with a covariate x , in general, as time passes, x and z will not remain uncorrelated. This gives rise to a number of difficulties and subtleties. One difficulty, first noted by Sheps and Menken (1973), is that unobserved heterogeneity can play havoc with attempts to make inferences about time or age dependence. A classic example is to consider a population which, when observed initially, is comprised of equal numbers of individuals from two groups, A and B . Suppose that in both groups, mortality is governed by a constant (exponential) hazard rate, with $r_A > r_B$, and suppose further that the analyst is ignorant of the existence of the two subgroups. Sheps and Menken observed that in this situation, the analyst will observe a monotonically declining rate with age, despite the fact that, by assumption, mortality does not vary with age in either group. This is an example of unobserved “frailty:” because members in A are frailer than those in B , they will die more quickly. As a consequence, the composition of the sample population will shift over time, moving from equal numbers of individuals from A and B to a population more heavily weighted towards those from B . Thus, what the analyst observes (technically, a distribution consisting of a mixture of the underlying r_A and r_B) will be mortality that is initially close to r_A but which will over time decline to r_B , as the sample composition of the population is increasingly selected against individuals from group A . See, for example, Trussell and Richards (1985) and Vaupel and Yashin (1985), with the latter providing informative examples showing how unobserved heterogeneity may affect conclusions about time variation in $r(t)$.

A second issue concerns potential biases in estimated covariate effects, an issue emphasized in a series of important and influential papers by Heckman and Singer (1982, 1984, 1985).

¹⁰A fixed-effects strategy, often used to analyze panel data on metric outcomes (Halaby forthcoming), was shown by Chamberlain (1985) to yield inconsistent estimates in an event history context; hence, researchers have concentrated attention on random effects models.

In reanalyses of unemployment data of Kiefer and Neumann (1981), employing a Weibull specification for duration dependence, they demonstrated the sensitivity of covariate estimates when unobservables were assumed to follow normal, log normal, and gamma distributions.¹¹ Their results, coupled with a series of Monte Carlo studies, led them to advocate an alternative approach to modeling unobserved heterogeneity using a discrete mixing distribution with finite points of support. One difficulty with this approach is that while it proceeds nonparametrically with respect to the distribution of the unobservables, it nevertheless requires strong parametric assumptions concerning the distribution of event times (Trussell and Richards 1985). Yet another possibility would be to proceed with reasonably flexible specifications for both the distribution of the unobservables (e.g., using a discrete mixing distribution) and of the event times (e.g., using the piecewise constant or spline specifications in [10] or [11]). For theoretical results concerning such an approach, see Elbers and Ridder (1982) and Heckman and Singer (1984).

Another approach to handling unobserved heterogeneity is due to Lillard and colleagues (see, e.g., Lillard 1993; Lillard and Waite 1993; Lillard, Brien, and Waite 1995; Upchurch, Lillard, and Panis 2001). It has long been recognized that when multiple transitions are observed for sample individuals, this added information would permit identification of additional distributional aspects of the unobservables. Lillard and colleagues have built upon this insight to estimate models akin to simultaneous equation models for metric outcomes in which unobservables may be correlated across transitions. Consider extending the proportional specification in (17) for multiple origin and destination states by adding an error term e_{jk} , where e_{jk} is assumed to capture unobserved heterogeneity specific to the transition from origin state j to destination state k :

$$\log r_{jk}[t|\mathbf{x}(t)] = \log q_{jk}(t) + \mathbf{x}_{jk}(t) + e_{jk} . \quad (18)$$

Let $m = 1, \dots, M$ index the full set of transitions $j \rightarrow k$ and let the e_{jk} be assumed normally

¹¹Heckman and Singer (1982) motivated their use of a Weibull distribution on search theoretical grounds. In subsequent work (Heckman and Singer 1985), they proposed an alternative Box-Cox-type parameterization for time dependence that has as special cases the exponential, Weibull, and Gompertz models.

distributed with mean 0 and covariance matrix

$$\text{cov}(e_m, e_p) = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & \cdots & \sigma_{1M} \\ \sigma_{21} & \sigma_2^2 & \cdots & \sigma_{2M} \\ \vdots & \vdots & \vdots & \vdots \\ \sigma_{M1} & \sigma_{M2} & \cdots & \sigma_M^2 \end{pmatrix}, \quad (19)$$

with the unobservables for two transitions, m and m' , correlated according to:

$$\rho_{mm'} = \frac{\sigma_{mm'}}{\sigma_m \sigma_{m'}}. \quad (20)$$

Thus, the models in (18)–(20) link different life course domains using observed *and* unobserved attributes of individuals and by allowing the effects of such unobservables to be correlated across life domains. Lillard, Brien, and Waite (1995) used such a modeling strategy to address whether the observed higher rate of divorce among those who have cohabited prior to marriage is an artifact of unobservables that differentially select couples into cohabitation. Thus, an unusual strength of these models is that they let researchers address endogenous selection. Nevertheless, it can be difficult to achieve identification of the many parameters in the variance-covariance matrix in (19), with identification achieved in most empirical work to date by imposing exclusion restrictions or by exploiting data on repeated events. Alternatively, one might achieve identification structurally via instrumental variables, although finding adequate instruments can be difficult in practice.

It is nevertheless important to emphasize that all models for unobserved heterogeneity proposed to date assume that what is unobserved does not vary with time and can be proxied by a one-dimensional correction term. Often, these assumptions are plausible; this is especially true for economic models of behavior, where it is commonly assumed that behavior is determined by a parsimonious set of influences, including those not observed by the researcher. Moreover, in many economic models of behavior, individuals are assumed to act rationally, optimizing over the life span; under these assumptions, characteristics such as permanent income and long-term horizons play an important role. By contrast, a recurrent theme among many life course researchers, particularly those drawn from sociology and demography, concerns the relatively fluid nature of the

social circumstances of individuals. If so, then extant models for unobserved heterogeneity address only one aspect of is unobserved—those unobserved characteristics of individuals that are fixed and unchanging—but not other aspects of what may be unobserved—unobserved characteristics of individuals that may be more fluid in nature.

Coupled processes.

To this point, the models I have considered focus on a single individual's life course transitions, but as noted earlier, life course theorists have often posited linkages across the life courses of multiple individuals—for example, that the events, behaviors, and outcomes for one member of a married couple might have profound effects on the events, behaviors, and outcomes for the person's spouse or partner. If such events, behaviors, or outcomes for both members of the couple are observed, if such covariates for a spouse or partner are exogenous to the outcome of the other, and if, conditional on these covariates, the processes for members of the couple can be assumed independent, then the problem reduces to the usual one of modeling effects of covariates on a transition of interest, with the set of covariates now expanded to include observed characteristics of a spouse or partner.

What is usually deemed implausible is the assumption of conditional independence; said another way, we often suspect that the set of observed covariates do not exhaust the set of what we would wish to observe theoretically and, in particular, that certain key characteristics for a couple are unobserved. A classic example concerns the problem of modeling the mortality of married couples, where the assumption of conditional independence is suspect if researchers do not observe aspects of diet or health-related behaviors that might affect both members of the couple, cumulative but unobserved health insults that are reflective of a couple's physical or social environment, or mate selection on unobserved characteristics that would tend to make couples more similar on health outcomes than two randomly chosen individuals in the population.

The specifications in (18)–(20) can, in principle, be adapted to cover some, but not all,

of these cases. Suppose, for example, that some unobserved characteristic of couples generates a positive correlation between their mortality experience relative to two individuals drawn at random from the same population. Concretely, let T_1 and T_2 denote the mortality of husbands and wives, respectively; then a model analogous to (18)–(20) for the mortality for couple i can be written as:

$$\begin{aligned}\log r_{i1}[t|\mathbf{x}_{i1}(t)] &= \log q_1(t) + \mathbf{x}_{i1}(t) + e_{i1} \\ \log r_{i2}[t|\mathbf{x}_{i2}(t)] &= \log q_2(t) + \mathbf{x}_{i2}(t) + e_{i2}\end{aligned}\tag{21}$$

with $\mathbf{x}_{i1}(t)$, $\mathbf{x}_{i2}(t)$, e_{i1} , and e_{i2} denoting observed covariates and unobserved components for the husband and wife in couple i , and where

$$\rho = \frac{\sigma_{12}}{\sigma_1^2 \sigma_2^2}$$

gives the correlation between members of a couple in age-specific mortality risks. As before, this model incorporates strong behavioral assumptions, in particular, that the unobservables e_{i1} and e_{i2} do not vary over time; note, in particular, that such an assumption would not cover the case in which mortality is affected by cumulative health insults as shared by a couple. In addition, identification of model parameters in (21) can be difficult when the data available to the researcher do not contain instruments that would, for example, plausibly affect the mortality of husbands, but not wives. An alternative approach is to model the *joint* distribution, $f(t_1, t_2)$, for the two event times, which yields so-called bivariate survivor models. In practice, this approach also has proven difficult to implement, in part because researchers often have little guidance for specifying the parametric form of the resulting two-dimensional baseline hazard function $q(t_1, t_2)$. Mare and Palloni (1988), Mare (1994), and Poetter (2001) provide a empirical examples and comprehensive discussions of these and other issues.

Conclusion.

Good methods often help sharpen theory—the translation of theoretical ideas stated verbally into testable propositions linked to data very often provides greater theoretical insight, for example,

by revealing conceptual ambiguities or gaps in a theoretical formulation. In this chapter, I have reviewed some examples that typify how the interplay between life course theory and event history methods might yield this sorts of analytical insights. Examples include: consideration of the populations deemed to be at “risk” of a particular life course transitions; careful specification of the states and transitions between states characterizing a problem not excluding those transitions of theoretical importance even if they may be difficult to observe empirically; how one might operationalize notions of age grading in the timing of various adult transitions, including extensions to other temporal dimensions—not just age—that might plausibly govern such transitions; how one might distinguish, conceptually and theoretically, notions such as marital “retreat” vs. “delay” for successive birth cohorts; what might be meant by an assertion about linkages across domains in the life course for a given individual and similarly what might be meant by an assertion that couple processes are linked.

Another theme running throughout this chapter concerns the implicit tradeoffs between nonparametric, semiparametric, and parametric methods. In my own work, I have found it useful to begin with exploratory analyses that make heavy use of nonparametric techniques—because these methods make few parametric assumptions, they can often provide important indications for how one might formulate more parametric models, especially when choosing among different models for the baseline hazard. I have also found it useful to use, where possible, flexible parametric models, for example, those utilizing piecewise linear splines in the place of a simple linear (or other parametric) specification. Such techniques can be used to model nonlinear effects of covariates or nonproportional effects of covariates; piecewise linear splines also yield quite flexible models for the baseline hazard. Note also that such models can be used to obtain tests of standard assumptions—for example, assumptions concerning linearity or proportionality in the effect of a covariate.

The ability to relax such assumptions has led to increased interest in nonparametric and semiparametric methods in the statistical and econometric literatures (see, e.g., Härdle 1990;

Hastie and Tibshirani 1990, 1993). Thus, when sufficiently large sample sizes are available to the researcher, it is often possible to devise methods that rely more heavily on information contained within the available data, for example, using this information to guide the choice of an appropriate functional form (e.g., as opposed to linearity or proportionality). Conversely, important insights follow from those models in which it is not possible to relax maintained assumptions in these ways (Manski 1995)—in such cases, estimated parameters are identified on the basis of assumptions that cannot be tested, even when arbitrarily large amounts of data are available. It is worth noting that, to date, while the classic work of Heckman and Singer (1982, 1985) on models for unobserved heterogeneity provides a semiparametric framework for uncorrelated unobservables, researchers have not yet devised nonparametric or semiparametric alternatives to the multiple transition models with correlated unobservables such as those in Equations (18)–(20).

The intersection of these issue—of how formal methods may clarify theoretical ideas and how nonparametric techniques can shed light on model identification—are useful in thinking about a central assertion in the life course: that one cannot understand seemingly disparate life course domains in isolation or that the life courses of spouses or partners exert mutual influences on one another. For example, consider life course transitions defined by cross-classifying an individual's social statuses across two or more life course domains—Figure 1 presents such an example for the transition of women through a cross-classification of fertility and marital statuses. It can be shown that standard nonparametric methods will reproduce the observed distribution across persons and across time through these multiple transitions and statuses. This result implies that the notion of linked life course domains must lie at some deeper level than the distribution of individuals across statuses over time.

One possibility is that the researcher directly observes all relevant aspects of that which is hypothesized to drive such linkages across life course domains. If so, one can proceed in the usual way by incorporating these data as standard covariates in a hazard regression. More typically, however, researchers worry that one or more key covariates driving such linkages are unobserved;

if so, then such unobservables will, in general, induce correlations across the life course transitions observed for individuals. The models outlined in Equations (18)–(20), which were developed to address precisely this problem, are thus of great intrinsic substantive interest to life course, since they allow researcher to address linkages and to obtain point estimates of the correlation between transitions for a sample of individuals or for the transitions observed for members of a couple paired. But as noted above, it has been difficult to date to devise obvious nonparametric or semiparametric alternatives to such models; hence, it remains unclear the degree to which parameter estimates from such models are identified solely from untestable model assumptions. These statistical difficulties may, in turn, be reflective of gaps in theoretical accounts—for example, ambiguities or incompleteness by life course theorists in specifying the range of theoretical and behavioral mechanisms that might generate behavioral linkages across different parts of the life course or across individuals in a couple.

Appendix: Software

It is important to emphasize that because software for event history models continues to evolve rapidly, any survey of available software will become dated rapidly. This being said, several readily available software packages (e.g., SAS, SPSS, Stata, S, S-Plus, and R) have modules that permit estimation of most basic models, including nonparametric estimation of the survivor, integrated hazard, and (sometimes) hazard functions, and estimation of the Cox model. Many of these packages also have provisions that permit the user to define time-varying covariates; in such cases, these models can be used to obtain estimates of the piecewise constant baseline model in Equations (9) and (10) by the appropriate coding of time-varying dummy variables representing the appropriate time-intervals.

Some of these packages also permit the estimation of models with multiple origin and destination states. When this is not possible, one can obtain estimates for standard models (i.e., those that assume conditional independence in the transitions to the multiple destination states) using a two-step procedure: first, by separating the problem into each origin state, and second, for each origin state, estimating parameters for a given destination state k by treating as censored those transitions to the other destination states.¹² Allison (1995) provides a useful and comprehensive survey of hazard estimation using SAS.

There are also several packages that have been developed for estimation of less standardized event history models. An incomplete list of such packages includes aML (Lillard and Panis 2000), CTM (Yi, Honoré, and Walker 1987), RATE (Tuma 1979), and TDA (Blossfeld, Hamerle, and Mayer 1989), and a variety of supplemental libraries developed for S, S-Plus, and R (see, e.g., Loader 2000). Most (but not all) of these specialized packages also permit estimation of basic models, but they otherwise differ considerably in the model coverage; hence, it is difficult

¹²This follows because maximum likelihood estimation for the expression in (15) can be shown to be separable by origin; intuitively, this follows because the origin state j determines the set of individuals at risk of the destination states k . Note that such an estimation strategy cannot be used for models that weaken the conditional independence assumption, for example, for those models in which terms for unobserved heterogeneity affect the transitions to the multiple destination states.

to identify any one package as providing superior coverage relative to another. For example, aML and RATE permit estimation of the piecewise linear spline for the baseline hazard function; provisions for piecewise linear splines in these packages also permit estimation of nonlinear and nonproportional covariate effects. TDA provides an extensive array of more parametric baseline hazard functions such as the Gompertz, Weibull, log logistic, log Gaussian, and sickle models. Local likelihood models (Wu and Tuma 1990) are most easily estimated using supplemental libraries available in S, S-Plus, and R (Loader 2000). CTM and aML have the most comprehensive routines for estimating models for unobserved heterogeneity, with CTM permitting estimation of the models discussed by Heckman and Singer (1985) and aML the models discussed by Lillard (1993).

REFERENCES

- Aalen, Odd O. 1978. "Nonparametric Inference for a Family of Counting Processes." *Annals of Statistics* 6(4): 701–26.
- Abbott, Andrew. 2000. "Reply to Levine and Wu." *Sociological Methods and Research* 29(1): 65–76.
- Abbott, Andrew, and Angela Tsay. 2000. "Sequence Analysis and Optimal Matching Methods in Sociology: Review and Prospect." *Sociological Methods and Research* 29(1): 3–33.
- Allison, Paul D. 1985. *Event History Analysis: Regression for Longitudinal Event Data*. Beverly Hills: Sage.
- Allison, Paul D. 1995. *Survival Analysis Using the SAS® System: A Practical Guide*. Cary, N.C.: SAS Institute.
- Andersen, P. K. and R. D. Gill. 1982. "Cox's Regression Model for Counting Processes: A Large Sample Study." *Annals of Statistics* 10(4): 1100–20.
- Becker, Gary S. 1991. *A Treatise on the Family*. (Enlarged edition.) Cambridge, MA: Harvard University Press.
- Blau, Peter M. 1977. *Inequality and Heterogeneity: A Primitive Theory of Social Structure*. Free Press.
- Blau, Peter M., and Otis Dudley Duncan. 1967. *The American Occupational Structure*. New York: Wiley.
- Blossfeld, Hans-Peter, Alfred Hamerle, and Karl Ulrich Mayer. 1989. *Event History Analysis*. Hillsdale, NJ: Erlbaum.
- Bumpass, Larry L. 1990. "What's Happening to the Family: Interaction Between Demographic and Institutional Change." *Demography* 27(4): 483–98.
- Chamberlain, Gary. 1985. "Heterogeneity, Omitted Variable Bias, and Duration Dependence." Pp. 3–38 in James J. Heckman and Burton Singer (Eds.), *Longitudinal Analysis of Labor Market Data*. Cambridge: Cambridge University Press.
- Cox, D. R. 1972. "Regression Models and Life Tables" (with discussion). *Journal of the Royal Statistical Society* B34(2): 187–220.
- Cox, D. R. 1975. "Partial Likelihood." *Biometrika* 62(2): 269–76.
- Cox, D. R., and D. Oakes. 1984. *Analysis of Survival Data*. London: Chapman and Hall.
- Dechter, Aimée R. 1992. "The Effect of Women's Economic Independence on Union Dissolution." Working Paper No. 92-28, Center for Demography and Ecology, University of Wisconsin-Madison, Madison, WI.
- Elbers, C., and G. Ridder. 1982. "True and Spurious Duration Dependence: The Identifiability of the Proportional Hazard Model." *Review of Economic Studies* 49(3): 403–10.
- Elder, Glen H. Jr. 1998. "The Life Course as Developmental Theory." *Child Development* 69(1):1–12.
- Elder, Glen H. Jr. 1999. *Children of the Great Depression: Social Change in Life Experience*. 25th Anniversary Edition. Boulder, CO: Westview.

- Fleming, Thomas R., and David P. Harrington. 1991. *Counting Processes and Survival Analysis*. New York: Wiley.
- Gilder, George F. 1986. *Men and Marriage*. Gretna, LA: Pelican.
- Glenn, Norval. Forthcoming. "Distinguishing Age, Period, and Cohort Effects and Synthetic Strategies." To appear in Jeylan T. Mortimer and Michael J. Shanahan, (Eds.), *Handbook of the Life Course*. New York: Plenum.
- Halaby, Charles N. Forthcoming. "Panel Models for the Analysis of Growth and Change in Life Course Studies." To appear in Jeylan T. Mortimer and Michael J. Shanahan, (Eds.), *Handbook of the Life Course*. New York: Plenum.
- Härdle, Wolfgang. 1990. *Applied Nonparametric Regression*. Cambridge: Cambridge University Press.
- Hastie, Trevor, and Robert Tibshirani. 1993. "Varying-coefficient Models." (With discussion.) *Journal of the Royal Statistical Society B55(4):757–96*.
- Hastie, Trevor, and Robert Tibshirani. 1990. *Generalized Additive Models*. London: Chapman and Hall.
- Heckman, James J., and Burton Singer. 1982. "The Identification Problem in Econometric Models for Duration Data." Pp. 39–77 in W. Hildenbrand (Ed.), *Advances in Econometrics*. Cambridge: Cambridge University Press.
- Heckman, James J., and Burton Singer. 1984. "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data." *Econometrica* 52(2): 271–320.
- Heckman, James J., and Burton Singer. 1985. "Social Science Duration Analysis." Pp. 39–110 in James J. Heckman and Burton Singer (Eds.), *Longitudinal Analysis of Labor Market Data*. Cambridge: Cambridge University Press.
- Hogan, Dennis P., and Nan Marie Astone. 1986. "The Transition to Adulthood." *Annual Review of Sociology* 12: 109–30.
- Hougaard, Philip. 2000. *Analysis of Multivariate Survival Data*. New York: Springer-Verlag.
- Kaplan, E. L., and Paul Meier. 1958. "Nonparametric Estimation from Incomplete Observations." *Journal of the American Statistical Association* 53(282): 437–81.
- Kiefer, Nicholas M., and George R. Neumann. 1981. "Individual Effects in a Nonlinear Model: Explicit Treatment of Heterogeneity in the Empirical Job-Search Model." *Econometrica* 49(4): 965–79.
- Lancaster, Tony. 1990. *The Econometric Analysis of Transition Data*. Cambridge: Cambridge University Press.
- Levine, Joel. 2000. "But What Have You Done for Us Lately? Commentary on Abbott and Tsay." *Sociological Methods and Research* 29(1): 34–40.
- Lillard, Lee. 1993. "Simultaneous Equations for Hazards: Marriage Duration and Fertility Timing." *Journal of Econometrics* 56(1/2): 189–217.
- Lillard, Lee A., Michael J. Brien, and Linda J. Waite. 1995. "Premarital Cohabitation and Subsequent Marital Dissolution: A Matter of Self-Selection?" *Demography* 32(3): 437–57.
- Lillard, Lee A., and Linda J. Waite. 1993. "A Joint Model of Childbearing and Marital Disruption." *Demography* 30(4): 653–81.

- Manski, Charles F. 1995. *Identification Problems in the Social Sciences*. HUP
- Mare, Robert D. 1994. "Discrete-Time Bivariate Hazards with Unobserved Heterogeneity: An Indirect Contingency Table Approach." *Sociological Methodology* 24: 341–83.
- Mare, Robert D., and Alberto Palloni. 1988. "Couple Models for Socioeconomic Effects on Mortality Among Older Persons." CDE Working Paper No. 88-07. Center for Demography and Ecology, University of Wisconsin-Madison.
- Mayer, Karl Ulrich. 1988. "German Survivors of World War II: The Impact on the Life Course of the Collective Experience of Birth Cohorts." Pp. 229–46 in Matilda White Riley (Ed.), *Social Structures and Human Lives, Volume I*. Newbury Park: Sage.
- Mayer, Karl Ulrich, and Nancy Brandon Tuma. 1990. "Life Course Research and Event History Analysis: An Overview." Pp. 1–20 in Karl Ulrich Mayer and Nancy Brandon Tuma (Eds.), *Event History Analysis in Life Course Studies*. Madison, WI: University of Wisconsin Press.
- Meyer, John W. 1986. "The Self and the Life Course: Institutionalization and Its Effects." Pp. 199–216 in Aage B. Sørensen, Franz E. Weinert, and Lonnie R. Sherrod (Eds.), *Human Development and the Life Course: Multidisciplinary Perspectives*. Hillsdale, NJ: Erlbaum Associates.
- Modell, John, Frank Furstenberg, and Theodore Hershberg. 1976. "Social Change and Transitions to Adulthood in Historical Perspective." *Journal of Family History* 1(1): 7–32.
- Nelson, Wayne. 1972. "Theory and Applications of Hazard Plotting for Censored Failure Data." *Technometrics* 14(4): 945–66.
- Oppenheimer, Valerie Kincade. 1997. "Women's Employment and the Gain to Marriage: The Specialization and Trading Model." *Annual Review of Sociology* 23: 481–53.
- Popenoe, David. 1996. *Life Without Father*. New York: Free Press.
- Poetter, Ulrich. 2000. "A Multivariate Buckley-James Estimator." In T. Kollo, E-M. Tiit, and M. Srivastava (Eds.) *Multivariate Statistics: New Trends in Probability and Statistics, vol. 5*. Utrecht: VSP.
- Settersten, Richard A., Jr., and Karl Ulrich Mayer. 1997. "The Measurement of Age, Age Structuring, and the Life Course." *Annual Review of Sociology* 23:233–61.
- Sewell, William H., and Robert M. Hauser. 1975. *Education, Occupation, and Earnings: Education in the Early Career*. New York: Academic Press.
- Shanahan, Michael J. 1999. "Pathways to Adulthood in Changing Societies: Variability and Mechanisms in Life Course Perspective." Forthcoming, *Annual Review of Sociology*.
- Sheps, Mindel C., and Jane A. Menken. 1973. *Mathematical Models of Conception and Birth*. Chicago: University of Chicago Press.
- Spain, Daphne, and Suzanne M. Bianchi. 1996. *Balancing Act: Motherhood, Marriage, and Employment Among American Women*. New York: Russell Sage.
- Trussell, James, and Toni Richards. 1985. "Correcting for Unmeasured Heterogeneity in Hazard Models Using the Heckman-Singer Procedure." Pp. 242–76 in Nancy Brandon Tuma (Ed.), *Sociological Methodology 1985*. San Francisco: Jossey-Bass.
- Tsiatis, Anastasios. 1975. "A Nonidentifiability Aspect of the Problem of Competing Risks." *Proceedings of the National Academy of Sciences* 72(1): 20–22.
- Tuma, Nancy Brandon. 1979. *Invoking RATE*. (2nd edition.) Menlo Park: SRI International.

- Tuma, Nancy Brandon, and Michael T. Hannan. 1984. *Social Dynamics: Models and Methods*. Orlando: Academic Press.
- Upchurch, Dawn M., Lee A. Lillard, and Constantijn W. A. Panis. 2001. "The Impact of Nonmarital Childbearing on Subsequent Marital Formation and Dissolution." Pp. 344-80 in Lawrence L. Wu and Barbara Wolfe (Eds.), *Out of Wedlock: Causes and Consequences of Nonmarital Fertility*. New York: Russell Sage Foundation.
- Vaupel, James W., and Anatoli I. Yashin. 1985. "Heterogeneity's Ruses: Some Surprising Effects of Selection on Population Dynamics." *American Statistician* 39(3): 176-85.
- White, Harrison C., Scott A. Boorman, and Ronald L. Breiger. 1976. "Social Structure from Multiple Networks. I. Blockmodels of Roles and Positions." *American Journal of Sociology* 81(4): 730-80.
- Winsborough, Halliman H. 1980. "A Demographic Approach to the Life Cycle." Pp. 65-75 in Kurt W. Back (Ed.), *Life Course: Integrative Theories and Exemplary Populations, AAAS Symposium 41*. Boulder: Westview.
- Wu, Lawrence L. 1989. "Issues in Smoothing Empirical Hazards." Pp. 127-59 in Clifford C. Clogg (Ed.), *Sociological Methodology 1989*. Washington: American Sociological Association.
- Wu, Lawrence L. 1990. "Simple Graphical Goodness-of-fit Tests for Hazard Rate Models." Pp. 184-99 in Karl Ulrich Mayer and N. B. Tuma (Eds.), *Event History Analyses in Life Course Research*. Madison, WI: University of Wisconsin Press.
- Wu, Lawrence L. 1996. "Effects of Family Instability, Income, and Income Instability on the Risk of a Premarital Birth." *American Sociological Review* 61(3): 386-406.
- Wu, Lawrence L. 2000. "Some Comments on 'Sequence Analysis and Optimal Matching Methods in Sociology: Review and Prospect.'" *Sociological Methods and Research* 29(1): 41-64.
- Wu, Lawrence L., and Brian C. Martinson. 1993. "Family Structure and the Risk of a Premarital Birth." *American Sociological Review* 58(2): 210-32.
- Wu, Lawrence L., and Elizabeth Thomson. 2001. "Racial Differences in Family Change and Early Sexual Initiation." *Journal of Marriage and the Family* 63(3): 682-96.
- Wu, Lawrence L., and Nancy Brandon Tuma. 1990. "Local Hazard Models." Pp. 141-80 in Clifford C. Clogg (Ed.), *Sociological Methodology 1990*. Washington, D.C.: American Sociological Association.
- Xie, Yu. 1994. "Log-Multiplicative Models for Discrete-Time Discrete-Covariate Event-History Data." Pp. 301-340 in Peter V. Marsden (Ed.), *Sociological Methodology 1994*. Washington, D.C.: American Sociological Association.
- Yamaguchi, Kazuo. 1991. *Event History Analysis*. Newbury Park: Sage.
- Yi, Kei-Mu, Bo Honoré, and James Walker. 1987. "CTM: A Program for the Estimation and Testing of Continuous Time Multi-State Multi-Spell Models." NORC/ERC and University of Chicago.

Table 1. Comparison of covariate effects for Cox, 3-period exponential, and 3-period splined Gompertz models. Age at first marriage, white women, 1980 Current Population Survey.

	Cox	3-perod exponential	3-period Gompertz
<i>Years of schooling completed</i>			
0–11	.25*** (.01)	.18*** (.01)	.25*** (.01)
13–15	–.35*** (.02)	–.35*** (.02)	–.35*** (.02)
16+	–.71*** (.02)	–.68*** (.02)	–.70*** (.02)
<i>Year of birth</i>			
1930–1949	.48*** (.01)	.54*** (.01)	.48*** (.01)
1950+	.16*** (.01)	.16*** (.01)	.17*** (.01)

Note: Standard errors in parentheses. See text for additional details.

* $p < .01$ ** $p < .001$ *** $p < .0001$ (two-tailed test)

Table 2. Selected comparisons of model fit for the 3-period exponential and 3-period splined Gompertz models. Age at first marriage, white women, 1980 Current Population Survey.**Panel A:** Model description and statistics

Model	model for baseline hazard	covariates?	$\log \mathcal{L}$
1.	3-period constant	no	-214,452.8
2.	3-period constant	yes	-212,293.0
3.	3-period linear spline	no	-207,678.2
4.	3-period linear spline	yes	-205,490.1

Panel B: Model comparisons

comparison	test for	df	χ^2
1 vs. 2	adding covariates	5	4,319.6
3 vs. 4	adding covariates	5	4,376.2
1 vs. 3	3-period constant vs. linear spline	1	13,549.2
2 vs. 4	3-period constant vs. linear spline	1	13,605.8




Figure 2: Hypothetical examples of a proportional decline in the hazard rate of first marriage for three successive birth cohorts. Implications for $S(t)$, $H(t)$, $r(t)$, and $\log r(t)$.




Figure 3: Hypothetical example of a “pure” delay in the hazard rate of first marriage for three successive birth cohorts. Implications for $S(t)$, $H(t)$, $r(t)$, and $\log r(t)$.




Figure 4: Nonparametric estimates of $S(t)$, $H(t)$, and $r(t)$ for white and black women, 1980 Current Population Survey.

Center for Demography and Ecology
University of Wisconsin
1180 Observatory Drive Rm. 4412
Madison, WI 53706-1393
U.S.A.
608/262-2182
FAX 608/262-8400
comments to: lwu@ssc.wisc.edu
requests to: cdepubs@ssc.wisc.edu