CHAPTER 1

Event History and Social Science

Social scientists often examine events, for example, the occurrence of a militarized dispute, unemployment, or adoption. Events like these connote change or represent a transition from one state to another. Frequently, this concern with events is tied to an interest in the history preceding the event, for example, the number of years leading up to a war or the number of months a child is in foster care before adoption. History, thought of in this way, involves timing, and for many research questions, the timing of social change is at least as interesting as understanding the event culminating the history. Such questions naturally lend themselves to an examination of both the occurrence of an event and the history leading up to the event's occurrence.

The issues of timing and change are relevant for social science and bear on many hypotheses and theories with which social scientists regularly work. Such hypotheses and theories may have observable implications related to timing and change. Moreover, methods accounting for timing and change often naturally follow from hypotheses or theoretical expectations embedded in the research question. Understanding an "event history" entails a consideration of not only *if* something happens, but also *when* something happens. An event history is longitudinal and event history *analysis* typically involves the statistical examination of longitudinal data collected on a set of observations. While a wide variety of statistical models may be constructed for event history data, at the most basic level, all event history models have some common features.

The dependent variable measures the duration of time that units spend in a state before experiencing some event. Generally, a researcher knows when the observations enter the process and when the process ends (with the occurrence or nonoccurrence of some event). Analysts are frequently interested in the relationship between the length of the observed duration and independent variables, or *covariates*, of theoretical interest. A statistical model may then be constructed linking the dependent variable to the covariates. Inferences can be made regarding the influence of the covariates on the length of the duration and the occurrence (or nonoccurrence) of some event. In the remainder of this chapter, we point out why event history models are suitable to a wide range of issues dealt with by social scientists.

The Substantive Motivation for Event History Analysis

Many of the problems, hypotheses, and theories underlying social science research have, at their core, an implicit or explicit interest in the notions of timing and change. Even if a researcher does not explicitly think in terms of "duration," that is, how long something persists before it changes, many interesting problems in the social sciences have observable implications that *are* longitudinal. By thinking of problems in terms of the longitudinal implications embedded in them, a potentially richer understanding of the social process underlying the problems can be achieved. We point out how some of the common themes and concerns in social science analysis are directly relevant to an event history model.

An Implicit Interest in "Survival"

Event history models are often referred to as survival models. Indeed, the class of models discussed in this book have a wide variety of names: duration models, survival models, failure-time models, reliability models, and so forth. The nomenclature arises from the different kinds of applications for which these models have been employed. For example, in engineering research, industrial reliability testing has led to the consideration of duration models, as these models naturally address questions of interest: How long does a mechanical component work (or "survive") until it fails? Similarly, many of the kinds of questions asked in social science are implicitly related to a conception of survival: Why do political parties maintain control or fail to maintain control of a legislature? How do politicians keep their seats over time, even when political conditions are unfavorable to them? How does the number of children affect the duration of marriage? Why do military conflicts persist or fail to persist? Why does the "peace" between one set of countries last longer than the peace between another set? Why are some families seemingly stuck in poverty?

Each of these questions beckons the notion of survival. Political parties or politicians, in order to maintain control of a legislature or of a seat, must survive over a series of elections. The length of a military conflict, or conversely, the duration of peace between countries invokes the idea of survival. Disputes can "survive"; peace can persist. Most of the important theoretical issues social scientists grapple with have implications regarding survival. For example, one facet of democratic theory suggests that a functioning and healthy democracy should permit some semblance of citizen control over its elected officials. If politicians are habitually dependent upon voter support, political "survival" may be a natural implication of such fundamental concepts as representation and citizen control. Theorists studying patterns of democratization in developing countries may treat as an implication of democratic stability, the duration with which regimes persist. Theorists of political institutions may be interested in the relationship between institutional design and rules and the duration of time that politicians survive or stay within the institution. Criminologists study the effectiveness of alternative rehabilitation programs on whether and when someone returns to prison. Health economists study the duration of hospitalization. Examples abound; the point is, the concept of survival is pervasive in social science. Event history analysis, or *survival* analysis, is explicitly premised on the notion of survival.

An Implicit Interest in Risk

Just as many social science theories have implications relevant for survival, the concept of "risk" is equally prevalent and important in social science research. It is difficult to consider survival without also explicitly considering risk: given that a political party has maintained control of the legislature for three elections, what is the risk the party will fall subsequently? The notion of risk in political science, or in any scientific field for that matter, implies a conditional relationship with survival. As something persists—as it survives—what is the risk it will subsequently end? Usually, political science questions pertaining to survival and risk are asked in more complicated ways: given a change in electoral rules, what is the risk that a party which has held control of the legislature for three elections will fall in the subsequent election? This question, which invokes the notions of survival and risk, ties these concepts to some tangibly interesting factor: an observable change to the rules governing elections. The kinds of questions that relate survival and risk to important theoretical factors are replete in social science. We demonstrate throughout this book how this notion of risk is directly incorporated into an event history model.

Event History Analysis *Is* Comparative Analysis

Social science research often strives or purports to be comparative. Indeed, at some basic level, just about every empirical question asked in social science is comparative in nature: given variation across some theoretical attribute, how do cases vary on values of the dependent variable? This question is rudimentary, but comparative. Likewise, event history models are explicitly comparative statistical models. Unlike traditional time-series models, where a single entity is typically examined over time, event history data contains information on many observations (i.e. individuals, politicians, wars, conflicts, convicts, parties, patients, countries, and so on) over time. Inferences from event history models can be very powerful. Not only can some claims be made regarding survival and risk, but also, explicit comparative inferences can be made regarding differences across the cases.

For example, in studying the duration of time coalitional governments survive (King, Alt, Burns and Laver 1990; Warwick 1992; Diermeier and Stevenson 1999), event history methods permit researchers to make claims not just about the factors that precipitate the risk a government will fall (or "stop surviving"), but also, how differences across political systems are related to this risk. The inference is comparative in nature. Given that event history data are longitudinal and generated across many observations, comparative inferences are naturally obtained from *any* event history model. It is not unreasonable to claim that *all* event history models are comparative. Certainly, the analyst may choose not to think of his or her results in comparative terms, but this kind of interpretation is forthcoming from an event history model.

Growing Body of Longitudinal Data

Social scientists are amassing an ever-growing body of longitudinal data. In part, the accrual of this kind of data has simply been a function of consistent and long-term research programs and data collection efforts. The accumulation of this kind of data has also stemmed from the recognition among social scientists that much more powerful inferences and theory-testing is possible with longitudinal data. Concomitantly, in the social science methodology literature, research on methods for time series, panel data, time-series cross-section designs, event counts, and event history data has flourished in recent decades. An equally burgeoning literature has emerged regarding the application of these kinds of models to substantive social science problems.

The fact that a considerable body of longitudinal data exists, or can be readily constructed from extant data sources, helps to motivate the consideration of event history models. But data availability alone is not sufficient to motivate the use of a statistical model. Rather, given the readily increasing availability of longitudinal data, *coupled with* social scientists' interest in the notions of survival and risk, one is directly led to the consideration of event history models. It is one of the aims of this book to demonstrate that the event history model is a valuable method for addressing substantive social science problems. Application of the models herein should be a matter of course when one has a substantive problem that requires the comparative analysis of longitudinal data. Since these problems abound in the social sciences, the event history model is a natural model for analyses.

Conclusion

It is incontrovertible that the substantive focus of many social science research problems leads directly to the consideration of duration models. This will be one of the principal claims we make throughout this book. As a road map for what is to come, in Chapter 2, we discuss the foundational principles of event history analysis. We consider the structure of event history data and introduce some important issues that will be of concern to us throughout the book. Additionally, we present the mathematical "building blocks" upon which event history analysis rests and then proceed to explain why traditional statistical models are problematic in the face of duration data.[1]

In Chapters 3-6, we consider in detail, estimation and interpretation of duration models for so-called "single-spell" durations. Specifically, Chapter 3 deals with parametric duration models—that is, models where the underlying hazard rate is parameterized in terms of a distribution function. In Chapter 4 we present the nonparametric alternative to the models considered in Chapter 3. In particular, we discuss the critical innovations of Cox (1972). We will make the argument that in general, the Cox model in most applied settings *will be preferable* over its parametric alternatives. In Chapter 5, we consider so-called "discrete-time" event history models. As we will point out later, the "discrete-time" label can be misleading inasmuch as discrete-time models often are good approximations of otherwise continuous-time processes. Chapter 6 provides a discussion of model selection, including further elaboration on the issue of parametric versus nonparametric estimation.

In Chapters 7-10, we discuss complications that emerge in event history data. Specifically, Chapter 7 deals with the inclusion of time-varying covariates in duration models; Chapter 8 discusses the implementation of model diagnostics; Chapter 9 considers the issue of unobserved heterogeneity; and Chapter 10 considers models for multiple events. By "multiple events," we mean the case when events can occur repeatedly, or the case where different/multiple kinds of events can occur. In the last chapter, Chapter 11, we summarize our principal arguments, revisit some of the issues raised regarding the relevance of the event history framework for social science, and make some recommendations regarding the implementation of duration modeling methods.

[1]We do not present the Bayesian approach to event history analysis in this book. Interested readers should see Ibrahim, Chen, and Sinha 2001.

CHAPTER 2

The Logic of Event History Analysis

The lexicon of event history analysis stems from its historical roots in biostatistics. Terms like "death," "failure," and "termination" are natural for analyses of medical survival data, but may seem awkward for social science analysis. In the context of medical research, survival data usually consist of longitudinal records indicating the duration of time individuals survive until death (if death is observed). In analyzing survival data, medical researchers are commonly interested in how long subjects survive before they die. The "event" is death, while the duration of time leading up to the death, the "history," is the observed *survival time*. Analysts working with survival data may be interested in assessing the relationship between survival times and covariates of interest such as drug treatments.

Likewise, social scientists frequently work with "survival data," although such data are generally not thought of in terms of survival and death. Nevertheless, much of the data social scientists use are generated from the same kinds of processes producing survival data. Concepts like "survival," "risk," and "failure" are directly analogous to concepts with which social scientists work. Thus, the concept of survival and the notion of survival and failure times are useful starting points to motivate event history analysis.

Event history data are, as Petersen (1995) notes, generated from *failure-time processes*. A failure-time process consists of units (individuals, governments, countries, dyads) observed at some natural starting point or time-of-origin. At the time-of-origin, the units are in some state (for example, holding some elected office) and are observed over time. A unit, at any given point in the process, is "at risk" of experiencing some event. An event represents a change or transition from one state to another state (for example, losing office in an election). After the event is experienced, the unit is either no longer observed or is at risk of experiencing another kind of event (or returning to the previously occupied state). In some instances, units are not observed experiencing an event, i.e., no transition is made from one state to another. Such cases are treated as *censored*, because although the event may be experienced, subsequent history after the last observation point is unobserved.

Such a process may be called a "failure-time" process because units are observed at an initial point, survive for some length of time or spell, and then "fail" (i.e. experience the event of interest) or are censored. The notion of failure is, of course, directly relevant to medical research or mechanical engineering, for example, where units *really do* fail: patients die, generators seize. But in social science applications, the "failure" is more appropriately thought of as an event, where the event denotes a transition from one state to another. The "failure time" represents the duration of time units survive until they fail. In political science applications, the failure time is analogous to the duration of time a unit is in some political state until it experiences an event.

Event History Data Structures

To make the notion of a failure-time process tangible, we consider some typical event history data structures. The "event" in event history analysis represents a change or transition from one state or condition of interest to another. For example, if a researcher is interested in studying the duration of a militarized intervention, the intervention is observed from its origin time until it ends. The termination of the intervention is the event and it represents the transition from one state (being "in a dispute") to another state (being "out of a dispute").

The premise of event history analysis is to model *both* the duration of time spent in the initial state *and* the transition to a subsequent state, that is, the event. At a minimum, event history data contain information on when the units begin the process under study and information on the timing of the event's occurrence (if an event is observed within the span of the observation plan). The starting time is usually treated as some natural beginning point of a process or state. For example, if one is interested in studying legislative career paths in the U.S. Congress, a natural starting point for observing House members is after their first successful election.

Defining an appropriate starting point for a process is a theoretical issue, not a statistical issue. For example, in analyzing the timing until a strong challenger emerges against an incumbent in a legislative election, the researcher must specify what the natural starting time of an election is. Because there is no officially sanctioned "start time" for campaigns, the time-of-origin is determined by the researcher, using theoretical guidance.

Knowing the time-of-entry into the process is important because it provides a natural baseline from which to compare units and observe subsequent history. Each unit in an event history data set is presumed to enter the process at the same time. In terms of "calendar time," the time-of-origin may vary across observations, but in terms of "clock time," the starting point is generally treated as equivalent for all observations. Continuing the example of

TABLE 2.1: Example of Event History Data: Military Interventions

Intervention	Intervenor	Target	Duration	Contiguity[a]	C[b]
1	U.K.	Albania	1	0	0
46	El Salvador	Honduras	657	1	0
81	U.S.	Panama	274	0	1
184	Bulgaria	Greece	12	1	0
236	Taiwan	China	7456	1	0
278	Botswana	S. Africa	1097	1	0
332	Uganda	Kenya	409	1	1
467	Israel	Egypt	357	1	0
621	Malawi	Mozambique	631	1	1
672	India	Pakistan	173	1	0

[a] Intervenors and Targets separated by 150 miles of water or less are coded as contiguous; [b]C denotes"censored": disputes on-going as of 31 Dec. 1988 are treated as right-censored. Data are Pearson-Baumann Militarized Intervention Data (ICPSR 6035).

congressional careers, the entry time begins after the first election. However, this start time varies across members of Congress, e.g., some members may be first elected in 1992 and others in 1994. Although the calendar time differs (1992 versus 1994), the "clock" begins 'ticking" at the same relative position: after the first election to Congress.[1]

Aside from giving information on starting times, event history data also provide information on the occurrence of the event. After some initial start time, units are observed in a state until at a later date, an event is experienced and a transition from one state to another is observed. The length of time that passes between entry into the process and occurrence of the event is the survival time or the duration time. Since the event must occur after the starting time, survival times must be positive. Note that events may or may not be observed for all individuals in a study. By the time the last observation period ends, some units may still be surviving. Units not experiencing an event by the last observation period are known as "right censored" observations because history subsequent to the last observation is unobserved. To illustrate what event history data look like in the context of a social science data set, consider Table 2.1. Here, we reproduce a portion of the International Military Intervention data set constructed by Pearson and Baumann (1993; see also Pearson and Baumann 1989 and Pearson, Baumann, and Bardos 1989). In the table, 10 of the 520 interventions from the data set are displayed.

The first column of data is the intervention number, the second and third columns list the intervening state and the target state of the dyad. The fourth column denotes the duration (in days) that the intervention lasted. The fifth

[1]In practical applications of event history analysis, sometimes the time-of-origin is unobserved, as an observation enters the study already in the process. Such an observation is referred to as "left truncated" because all history prior to, or "left" of, the initial observation point is unobserved, or truncated. Later in this chapter we discuss the problem of left truncation in more detail.

column lists the values of a binary variable indicating whether the intervenor and target are geographically contiguous states. And the sixth column indicates whether the intervention is right-censored. Contained in the duration time are two important pieces of information. We know how long the intervention "survives" or lasts from its onset until its termination date. The duration time indicates both the length of the intervention and the time at which the event occurred. The event is the termination (or "failure") of the intervention. For right-censored interventions, note that the duration time only provides information on the intervention up to the last observation point; no information is revealed on when the intervention terminates. So for example, although cases 46 and 621 have similarly recorded duration times (657 and 631 days), we know that the dispute between El Salvador and Honduras (case 46) ended on the 657th day after its onset; the only thing we know about the dispute between Malawi and Mozambique (case 621) is that as of the last observation point, the dispute had survived for 631 days. The similarity in duration times is illusory: one intervention has ended, the other has not. It is not uncommon for event history data sets to contain numerous right-censored observations. Analyses that fail to distinguish uncensored and censored cases can produce misleading conclusions. This point is elaborated later in this chapter.

Finally, Table 2.1 includes information regarding a variable of interest, contiguity status. Typically, analysts are interested in studying the relationship between duration times and covariates. In the case of these data, it may be interesting to ask whether duration times of interventions vary according to contiguity status of the intervenor and target (cf. Goertz and Diehl 1992, Mansfield and Snyder 1995, Mitchell and Prins 1999). A model may be constructed treating the duration time as some function of contiguity status. From the model, a researcher could assess if this covariate is associated with longer or shorter duration times. Using the risk terminology discussed previously, one can assess if the "risk" of an intervention ending increases or decreases with contiguity status. The contiguity covariate in Table 2.1 is known as a "time-independent" covariate, as its values do not change within observations over time. However, researchers frequently will be interested in covariates that have values that change within observations over time. Such covariates may be referred to as "time-varying" covariates (TVCs). The use of TVCs helps to motivate event history analysis, but raises special problems as well. These problems are discussed in Chapter 7.

So far, our interest has implicitly centered on "single-spell" processes, or single event processes. In our discussion of militarized interventions, for example, we referred to the duration of *an* intervention and (conveniently) ignored the possibility that a state may become involved in multiple interventions. Likewise, in our example of House careers, we side-stepped the obvious issue that a career could end in a variety of ways, for example through electoral defeat or through retirement from office. In short, we have assumed there is a

TABLE 2.2: Disputes between Nicaragua and Costa Rica

Date of Dispute Onset	Date of Dispute Termination	Duration	Outcome
Dec. 11, 1948	Mar. 9, 1949	89	Stalemate
April 1, 1954	Feb. 24, 1955	330	Compromise
May 3, 1957	June 23, 1957	51	Stalemate
Oct. 10, 1977	Oct. 15, 1977	6	Stalemate
Sept. 12, 1978	Dec. 27, 1978	107	Stalemate
Sept. 28, 1983	Sept. 3, 1984	342	Stalemate
May 31, 1985	June 5, 1985	6	Stalemate
April 16, 1986	April 16, 1986	1	Stalemate
Sept. 2, 1987	Sept. 2, 1987	1	Stalemate

singular event of interest: the termination of a single spell. For many research questions, however, the focus on a single event (or transition) is limiting. Event histories are often complicated. The most general form of an event history can encompass multiple events of the same type, or multiple events of different types. Moreover, the complications that emerge in event history structures can directly influence the modeling strategy one chooses.

To illustrate some of the issues that emerge in complicated event history data structures, consider the data in Table 2.2. These data are taken from a large data set recording the occurrence and duration of militarized interstate disputes, or MIDs (see Jones, Bremer, and Singer 1996 and Maoz 1999 for details on the MID data). The MID data have been widely used in studies of international conflict because the structure of the data set permits one to examine, longitudinally, disputes that occur between pairs of countries, commonly referred to as "dyads." In Table 2.2, we give the event history data for disputes that have occurred between Nicaragua and Costa Rica in the post World War II era;[2] hence, the data in Table 2.2 provide a comprehensive account of the dispute history between this dyad. We see that in this era, the two countries were involved in 9 disputes varying in length from 1 day (the MID began and ended on the same day) to 342 days. The data help to illustrate several points regarding event history data structures.

First, note that there are *multiple* spells, or duration lengths, each corresponding to a separate MID. This implies that this dyad has *repeatedly* engaged in disputes with one another. Second, note also that there are *multiple* events, or outcomes, in the dispute history between Nicaragua and Costa Rica: one dispute ended in compromise, the other disputes ended in a stalemate. Third, note that there is an implied duration time *between* disputes. That is, between dispute spells, there are nine "peace" spells, corresponding to the length of time between the termination of a prior dispute and the onset of a subsequent dispute. In piecing together the full dispute history for this dyad, an event

[2]This dyad was arbitrarily chosen from the MID Data to illustrate event history structures.