CHAPTER 32

Event History Methods

Janet M. Box-Steffensmeier and Anand E. Sokhey

AN INTRODUCTION TO EVENT HISTORY ANALYSIS

Researchers are often interested in more than just the occurrence or non-occurrence of a political events; often the *timing* of events is of equal substantive importance, whether it is the dissolution of a government's cabinet (e.g., King et al. 1990; Warwick 1992; Diermeier and Stevenson 1999), the presence of international military disputes (Jones et al. 1996; Werner 2000; King and Zeng 2001), contributions by political action committees (Box-Steffensmeier et al. 2005), or as we will examine in this paper, when a voter makes up his/her mind in an election campaign. Examining when an event occurs provides additional information and may lead to new insights into the event and process under study. Event history – or survival analysis – is the tool of choice when political scientists find that the answer to "why" necessitates an answer to "when."[1]

At its base, event history involves the statistical analysis of data that is longitudinal in nature or that at least implies a longitudinal process. The dependent variable is the amount of time that an observation – whether a country, dyad, individual, etc. – spends in one state before entering another; in the case of a voter choosing a candidate in an election, it would be the amount of time that the individual spends making up his/her mind before deciding whom he/she is voting for (i.e., the amount of time before changing from the state of "undecided" to the state of "decided").

Such state changes are typically referred to as "failures" or "events," and depending on whether a discrete or continuous approach is taken, these can occur either anywhere in time (the continuous case), or only within observed intervals (the discrete case). In this chapter, we will focus on models that assume continuous event history processes, although discrete-time models remain a popular alternative (see Beck 1999 for a discussion).[2]

[1] We wish to distinguish our focus in this chapter – time duration modeling – from other recent approaches developed for analyzing "events data." For example, Schrodt and colleagues have been applying hidden Markov models in the realm of international relations; the idea is to compare sequences of discrete events to produce quantitative estimates that match "precedent-based reasoning" (for an excellent overview, please see Schrodt 2000).

[2] Researchers have – sometimes incorrectly – used a variety of techniques for analyzing discrete duration data, including conventional logit/probit (with or without added time dummy variables; for a discussion, see Beck et al. 1998), and transition techniques (which consist of separate models for separate transition processes). Duration data analyzed with logit/probit models need to account for duration dependence, which is typically done with splines.

K.T. Leicht and J.C. Jenkins (eds.), *Handbook of Politics: State and Society in Global Perspective*, Handbooks of Sociology and Social Research,

Of course, researchers may also deal with data processes in which there are multiple failures (i.e., repeated events), or multiple spells (i.e., periods during which a subject is at risk of failing), or both. These additional data complications are straightforward to address and are critical to making correct inferences.

Since event history is concerned with the timing of change, it makes sense that analysis begins by conceptualizing survival times as a positive random variable, T, with a distribution function:

$$F(t) = \int_0^t f(u)d(u) = \Pr(T \leq t) \tag{1}$$

Differentiating $F(t)$ yields the probability density function $f(t)$,

$$f(t) = dF(t) / d(t) \tag{2}$$[3]

which like $F(t)$, characterizes the failure times. In turn, the survivor function, $S(t)$,

$$S(t) = 1 - F(t) = \Pr(T \geq t) \tag{3}$$

denotes the probability that a survival time T is equal to or greater than some time t. Pairing these two functions provides the *hazard rate*, $h(t)$,

$$h(t) = f(t) / S(t) \tag{4}$$

which captures the relationship between the density of failure times, $f(t)$, and the survivor function, $S(t)$. The concept of risk is at the heart of event history analysis, and the hazard rate is intimately tied to this idea – the hazard describes the rate at which observations fail by time t, given that they have survived up until t. Social scientists are often interested in understanding how this – the risk of an event – changes in response to the values of various independent variables or covariates.

THE STATISTICAL MOTIVATIONS FOR EVENT HISTORY ANALYSIS: DEALING WITH DURATION DEPENDENCE

In ordinary least squares (OLS) regression, the residuals (ε_i) are assumed to follow a normal distribution. Thus, if we were to model an event history process using such an estimation procedure, the time to an event – conditional on our covariates – would also be assumed to follow such a distribution. However, in thinking about real world failure-time processes, such an assumption would be both hard to justify (as the distributions of such times are often asymmetrical) and would often lead us to incorrect inferences (as OLS is not robust to such deviations) (Cleves et al. 2004).

Other primary statistical motivations for using the event history approach include censored data and time-varying covariates; OLS is an improper technique for modeling failure-time processes because of its inability to deal with these issues. Censoring occurs when an observation's full history is not observed, and event history analysis is specifically designed to account for censored data via the calculation of the hazard rate. For example, in studying the duration of an international military dispute, the dispute may be ongoing at the end-time of the analysis – that is, it has not ended, and thus the dispute is *right censored. Left-truncation* occurs when some observations have experienced an event before the beginning of the study;

[3] In the discrete case, the probability mass function for a discrete random variable is $f(t) = \Pr(T = t_i)$.

it can also be considered a censoring problem as data are not-observed, only in this case the nonobservation occurs prior to the start of the study.

Time-varying covariates are also readily handled in an event history analysis. Allowing the value of the covariates to change over time is important in order to properly assess hypotheses. For example, time-varying covariates are needed to gauge the impact of war chests on whether a challenger enters an electoral race – war chests need to be measured over the course of the election cycle as simply measuring them at one time point (whether at the beginning, middle, or end of the cycle) would be woefully inadequate (For a more in-depth discussion, see Box-Steffensmeier and Jones 2004: Chap. 2.).

Parametric Modeling

Parametric event history models improve upon OLS by directly modeling the duration dependence in the data using more appropriate distributional forms. For example, if we thought the "risk" of an individual making a decision among presidential candidates was *constant* over the course of the electoral campaign, specifying an exponential distribution for the time dependency would be the right choice as it characterizes the baseline hazard as flat; if we thought the risk of making a decision was monotonically increasing (or decreasing) over time, the Weibull, which nests the exponential, might be appropriate.[4] Other parametric models such as the log-logistic can offer the researcher a bit more flexibility in that they allow the specification of nonmonotonic hazard rates.

All such parametric models are estimated through maximum likelihood, with the likelihood function being expressed in terms of the density of whatever distribution one has chosen; most parametric models can be run fairly easily in popular software packages such as *R* and *Stata*. Best practice demands that the choice of distributional forms always be guided by theory, though as Blossfeld and Rohwer (2002) note, social scientists rarely have theory sufficient to justify a particular parametric choice. Further, the choice of parameterization is an important one, for different distributional assumptions can produce markedly different results (we will further address both of these points in the next section). To revisit the vote decision example, simply assuming that the "risk" of an individual's decision increases monotonically (e.g., as a Weibull) as a function of the approach of election day may be unwise; electoral politics research – and conventional political wisdom – would suggest that the hazard rate may be nonmonotonic (e.g., as a log-normal), rising and falling to reflect the major milestones of the campaign, including the parties' conventions and the presidential debates. If parametric models are used, extra testing is needed to determine if the appropriate parametric distribution has been chosen. As in regular maximum likelihood analysis, the fit of parametrically nested models may be compared using a likelihood ratio (LR) test. The fit of parametric, non-nested models can be compared by using the Akaike information criterion (AIC).

Semi-Parametric Modeling: The Cox Proportional Hazards Model

While the previous set of models makes distributional assumptions about the nature of the time dependency in the data, the Cox model (Cox 1972, 1975) leaves this, the baseline hazard, unspecified. While parametric models and the Cox model are both perfectly acceptable ways

[4] These models (along with others, such as the Gompertz) assume *proportional hazards*, which must be tested for during implementation. We define and discuss this model property in our discussion of the Cox semi-parametric model.

to proceed with event history estimation (both are widely used across numerous scholarly fields), the Cox model offers some key advantages for social scientists. A central benefit of the Cox approach is that it allows researchers to avoid the testing of various parametric assumptions *by allowing them to avoid having to make assumptions about the nature of the duration times in the first place* – assumptions which may be poorly informed, which may lead to incorrect inferences, and which are often of secondary importance to the relationship between the outcome variable and the set of covariates under consideration (Box-Steffensmeier and Jones 2004).[5] Accordingly, we consider the Cox to be a more straightforward alternative to parameterization techniques.

In the Cox model, the hazard rate for the ith individual is specified as:

$$h_i(t) = h_0(t)\exp(\beta'\mathrm{X}) \tag{5}$$

where $h_0(t)$ is the baseline hazard function, and $\beta'\mathrm{X}$ are the covariates and regression parameters; looking at the equation, we can see that changes in the baseline hazard are solely a function of the covariates and are a multiple of the baseline. Thus, like some of the aforementioned parametric models (e.g., the exponential and the Weibull), the Cox model also adheres to the *proportional hazards property* (hence it is sometimes called "the Cox proportional hazards model") – which means that this proportional change in the baseline is assumed to be fixed across time. Like all modeling assumptions, the proportional hazards assumption should always be tested for violations, and we will demonstrate common diagnostics and corrections in the example presented below.

Unlike the aforementioned models, however, the Cox proportional hazards model is estimated through *partial* maximum likelihood (i.e., not full MLE), so named because only part of the information available in the data is used in the estimation. Under this method, it is assumed that the intervals between events provide no information about the relationship between the covariates of interest and the baseline hazard. Rather, it is the ordered failure times that contribute information to the partial likelihood function – time matters to the extent that it gives order to the failure times (Cleves et al. 2004: 5; Collett 1994). It is this breakthrough that provides the tradeoff which allows the parametric assumptions to be relaxed.

To derive the partial likelihood function, we begin with the conditional probability of a failure at time t_i, given the number of cases that are in the "risk set" – that is, the number of cases that are at risk of failure at t_i. Equation 6 denotes the probability that the jth case will fail at time T_i, given the number of cases that are at risk at time t_i (defined by $R(t_i)$) (while summing over all individuals in the risk set).

$$\Pr(t_j = T_i \,|\, R(t_i)) = \frac{e^{\beta' x_i}}{\sum_{j \in R(t_i)} e^{\beta' x_j}} \tag{6}$$

Taking the product of the conditional probabilities produces the partial likelihood function (which is often logged before being maximized):

$$L_p = \prod_{i=1}^{K} \left[\frac{e^{\beta' x_i}}{\sum_{j \in R(t_i)} e^{\beta' x_j}} \right] \tag{7}$$

---

[5] Discrete and continuous time approaches are both acceptable ways to proceed with event history estimation. However, the continuous time approach – which we discuss here – is more straightforward. By using a continuous time approach, one does not have to fit and test for an appropriate link function to account for the duration dependence. However, see Beck et al. (1998) and Beck (1999) who argue that discrete time approaches are more straightforward to interpret due to researchers' familiarity with discrete time (i.e., logit and probit) models.

Given that the Cox model's partial likelihood function is based solely on the ordered failures in the data, estimation originally could not take place in the presence of "ties," or coterminous events. However, in the last couple of decades, approximation and computing advances concerning the risk set have solved this problem. The issue of ties is relevant for any continuous time model, but several methods exist for dealing with this problem, including the Breslow, Efron, and Exact Discrete methods. Indeed, another advantage of the Cox model over parametric models is its ability to deal with data that is heavily "tied" (Box-Steffensmeier and Jones 2004; Golub and Collett 2002).

EXTENSIONS TO THE COX SEMI-PARAMETRIC MODEL

Useful and important extensions to the basic Cox model include approaches for dealing with multiple events and unobserved heterogeneity – these include shared frailty (multilevel models) and individual frailty models. The flexibility of the Cox model to account for such unique data aspects – in addition to the extensive diagnostics available – has contributed to the popularity of the approach.

Multiple events can be unordered or ordered; unordered events are often referred to as *competing risks*, and ordered events as *repeated events*. Competing risks models allow the researcher to incorporate additional information about the data and to test more specific hypotheses. For example, we might be interested in not only whether or not a member of a legislature leaves office, but *how* the member leaves office – by retirement, scandal, defeat in the primary, defeat in the general election, or to run for higher office (see Box-Steffensmeier and Jones 1997); the nature of the event is important information, and we expect the effect of the covariates to vary based on these different types of events. Furthermore, ignoring this information could lead to incorrect inferences – the effects may be the opposite across the different types of events, and this would be missed if the researcher were to collapse all types into only one summary event.

Repeated events occur in a specific order, and taking into account this sequencing information – rather than treating all the events as independent – is likely to be important. For example, the hazard rate may vary or the covariate effects may differ for a child who has been placed in foster care for the fifth time versus for the first time (Box-Steffensmeier et al. 2008).

Cox models may also be extended to account for unmeasured, unmeasurable, or unknown sources of heterogeneity, and these statistical dependencies can be accounted for via shared or individual frailty models.[6] Therneau and Grambusch (2000) define a frailty as a continuous variable that describes excess risk for distinct categories such as individuals, families, countries, or regions. The idea is that observations have different frailties, and that those who are the most "frail" will experience the event first (2000: 231). Dependencies arise for a variety of reasons, including spatial location, such as observations being from the same legislative

[6] Another popular strategy for dealing with unobserved heterogeneity is through split-population models (Schmidt and Witte 1988), which relax the assumption that eventually all observations will experience the event of interest. With origins in the biostatistics literature (Boag 1949), these models split the sample into two groups: one that has some risk of experiencing an event and one that has essentially zero risk; overly high failure rates are avoided by adjusting the information that is contributed to the likelihood function by the low-risk "population" (for political science applications, see Box-Steffensmeier et al. 2005; Clark and Regan 2003).

district, state, country, or region. It is also worth noting that levels may be defined by distance rather than by a fixed region, such as the "Middle East." That is, all countries within 500 miles of each others' capitals (e.g., Qatar, UAE, and Oman) may be defined as having a shared frailty; in the case of distance, the frailties may overlap (see Banerjee et al. 2003).

Garibotti et al. (2006) point out that the shared frailty model is attractive, because it explicitly acknowledges the potential role of unobserved factors that affect the duration of the event being studied. They also note that it assumes that unobservable characteristics are perfectly shared with others in the specified group (such as the family, state, or school), and that unobserved factors that are not shared are not considered. In contrast, correlated frailty models allow for individual-level frailties that can be correlated across the individuals (or more generally, the observations) within a group. The shared frailty (or multilevel/random effects) Cox model is another useful extension that should be applicable across the social sciences.

Finally, the conditional frailty model is another Cox extension designed to account for the presence of both repeated events and heterogeneity through stratification and random effects. Stratification by event number, i.e., first, second, third, etc., occurrence of the event, provides the flexibility of varying baseline hazards to allow for event dependence, and the addition of a frailty term captures unmeasured variation in the dependent variable. Allowing for the possibility of event dependence *and* heterogeneity provides additional modeling flexibility (see Box-Steffensmeier and De Boef 2006; Box-Steffensmeier et al. 2007).

WHEN DO VOTERS MAKE UP THEIR MINDS?

Data: The 2004 American National Election Study

Having introduced the Cox model and some of its popular extensions, we now proceed to our example analysis. Using data from the 2004 American National Election Study, we analyze the timing of when voters decided which candidate to support in the 2004 presidential election.[7] The specific wording of the timing question is as follows: "How long before the election did you decide that you were going to vote the way you did?" The item immediately preceding this one on the survey instrument asked who the respondent voted for in the presidential election.

The general categories and distribution of responses for this question are provided in Table 32.1. Looking at the table, we note that 33% of respondents stated that they knew "all along" how they would vote. At the other end of the spectrum, over 15% reported deciding within the last 2 weeks of the campaign. We code the dependent variable in days, where day 1 indicates the earliest deciders and day 252 – election day – indicates the latest deciders.

To the best of our knowledge, no one has looked at the timing of the voting decision in quite this way.[8] We include a number of covariates to explain the timing of one's decision, and have divided these variables into three basic categories: personal political characteristics, factors related to political engagement, and demographic controls. The personal political characteristics include strength of partisan identification, strength of ideology, and disapproval of the president. We expect those who decide early in the election cycle to be strong partisans

[7] The 2004 American National Election Study is available through the Inter-University Consortium for Political and Social Research (ICPSR); http://www.icpsr.umich.edu/, accessed August 15, 2007.

[8] McClurg (2006) examines how social networks (and other factors) influence the timing of the decision to vote for a specific candidate, but does not conduct an event history analysis.

Table 32.1. When Voters Made Up Their Minds in the 2004 American Presidential Election

	Frequency	Percentage	Campaign days "survived"
"Knew all along/always/from the first/9 months or more"	276	33.62	1
"During/after the primaries/5–8 months before"	93	11.33	42
"Before the conventions/early on"	85	10.35	84
"At the time of the Democratic convention (7/26–7/29/04) /3–4 months"	50	6.09	154
"At the time of the Republican convention (8/30–9/2/04) /2–3 months"	71	8.65	182
"after the conventions/during the campaign/September/ a couple of months"	47	5.72	196
"5–7 weeks before"	8	0.97	217
"1 month/October/after the debates/several weeks"	66	8.04	231
" ~2 weeks/10 days before"	51	6.21	242
"in the last days/a week/less than a week"	56	6.82	247
"on election day"	18	2.19	252
Total	821	100.00	

Note: Respondents were asked the following question: "How long before the election did you decide that you were going to vote the way you did?" The question immediately preceding asks who the respondent voted for in the presidential election.
Source: The 2004 American National Election Studies.

and/or ideologues – their partisanship and ideology will serve as the strongest possible "cue" (Conover and Feldman 1981), and thus they will be among the first to know who they will support in the election.[9] In the spirit of retrospective evaluation (Fiorina 1981; Key 1966), we include a dummy variable for presidential disapproval to test whether voters who disapprove of the job the president has done will decide earlier in the campaign to vote against the president's party.

The political engagement variables include previous voting participation (in the 2000 election), the frequency of the individual's political conversations (with family, friends and peers), the respondent's level of political interest, and the respondent's level of political knowledge. For the first of these covariates, we test whether previous participatory experience leads a voter to an earlier decision on whom to support in the election. We expect that experienced voters may reach a decision sooner, reflecting greater political awareness (Zaller 1992) or perhaps political sophistication (Luskin 1987).

Regarding the second factor, we note that conversation serves to provide voters with information (Downs 1957; Huckfeldt and Sprague 1995), and thus posit that more frequent political discussion may provide an individual with more information earlier in an election (especially before the campaign is in the "home stretch," and all individuals are exposed to more electoral stimuli). In turn, we expect that such early "doses" of information may cause an earlier crystallization of opinion. We posit that political interest and knowledge (Delli Carpini and Keeter 1996) work in largely similar ways: individuals with greater interest seek out more information and are thus more likely to have better-formed opinions; those with higher stocks of political knowledge (i.e., those who are more familiar with government) are more likely to have stronger preferences (which again translates into earlier decisions).

[9] Relatedly, we might expect that stronger partisans and ideologues would be less ambivalent (Zaller 1992) about the presidential contest, which would make them more likely to come to a decision earlier (McClurg 2006).

Finally, we include a number of demographic controls to test whether there are differences by age, gender, education, income, and race. However, we do not expect such differences to emerge, as there is no theory that we are aware of to suggest – for example – that women decide on a candidate sooner or later than men.

In the model, we include a shared frailty term to account for the multilevel nature of the data – that is, voters (i.e., the lower-level units) are nested within state electoral contexts (i.e., the higher-level units). Because of the United States' electoral college and the winner take all system in the states, American presidential campaigns are strategic and state focused. Thus, we might expect that individuals in some states would be more likely to make up their mind earlier than individuals in other states. A shared frailty model estimates a random parameter to account for the unmeasured factors that make individuals in certain states more "frail" than individuals in other states when it comes to the timing of the decision – this makes sense as we have little "level-2" data in the American National Election Study.[10]

Results

Researchers should always conduct modeling diagnostic tests, and thus we begin by examining a few common procedures as they pertain to our models: the link test and test(s) of the proportional hazards assumption. A link test can be used to evaluate general model specification. The intuition behind the test, which can be applied beyond event history models, is to evaluate the specification of the model by testing an alternative specification; this specification is based upon a re-estimation which uses the transformation of linear predictors from the model being scrutinized. If the transformation is statistically indistinguishable from zero when included in the model with the linear predictions, then the model is well-specified; if not, the model has problems that require further inquiry. In the case at hand, for both models, a link test does not reveal any concerns, as the p-value is not statistically significant. For the shared frailty model, the positive coefficient on $\hat{y}^2$ (0.024) has a statistically insignificant *z*-score of 0.13, with a *p*-value of 0.90.

Since the Cox model belongs to the class of survival models that relies on the assumption that the covariates' effects on the hazard rate are proportional over time, one area of concern is the possibility that the proportional hazards assumption does not hold for one or more of the covariates. Using the Schoenfeld residuals (called during estimation), we can employ a number of diagnostic tests – both graphical and statistical – to determine whether there are any violations of the proportional hazards assumption in the model. We use the straightforward Grambusch and Therneau global test (Grambusch and Therneau 1994) for the models, as well as Harrell's rho (Harrell 1986) for individual covariates. These statistical tests avoid the subjectivity inherent in graphical tests that require trying to decide whether residuals fall in consistent, discernable patterns, or are randomly dispersed. We should emphasize, however, that the graphical tests *are* still useful, particularly when trying to determine which function

[10] When estimating a frailty model, a distribution must be specified for the random effects. Though the gamma is the typical – sometimes only choice in certain statistical packages (and what we present in the tables) – some scholars have criticized the often "atheoretical" choice of distributional forms, and others have noted that different assumptions can greatly alter the parameter estimates obtained (Blossfeld and Rohwer 2002). Frailty models are an active and fast-developing area of research, including semi-parametric approaches for the estimation of the frailty parameter (e.g., Andersen et al. 1999; Li and Lin 2003).

of time to use as an interaction with offending covariates when addressing violations. Our tests here show that there is not a problem with the proportional hazards assumption in either model – in the case of the frailty model, we see that the Global Proportional Hazards Test is not statistically significant with a χ^2 value of 9.6 and Prob > χ^2 0.65.

Looking at the results of the model (see Table 32.2), we find that some personal political characteristics, as well as some political engagement covariates, have a statistically significant effect on when voters decide which presidential candidate to support. Specifically, voters with stronger ideologies and stronger partisan identification decide earlier; those who voted in 2000 and those with greater political interest are also predicted to decide earlier.[11] For hazard ratios (presented in the second column of Table 32.2), estimates below one indicate a lower risk of experiencing the event – in this case making a decision – and estimates above one indicate a higher risk of experiencing the event, all else equal; the largest hazard ratio is for political interest. Moving down the column for the shared frailty model, we see that the estimated size of the random effect (akin to a random effect in a multilevel model) is 0.021, which is small,

Table 32.2. Predicting the Timing of Voting Decisions, 2004 (Cox Prop. Hazards Estimates; Multilevel (Shared Frailty) Model)[a]

Covariates	Coefficient	Hazard ratio	S.E.
Personal political characteristics			
Strength of ideology	0.21	1.24	0.05***
Strength of partisanship	0.24	1.27	0.05***
Disapprove of the President	−0.01	0.99	0.09
Political engagement			
Frequency of political discussion	0.02	1.02	0.02
Voted in 2000	0.24	1.27	0.14*
Political interest	0.35	1.42	0.09***
Political knowledge	0.07	1.07	0.05
Demographics			
Age	−0.00	1.00	0.00
Gender	0.04	1.05	0.09
Education	−0.01	0.99	0.02
Income	0.00	1.00	0.01
Race (African American)	0.24	1.28	0.16
Random effect (Shared frailty term θ)	Variance: 0.021	Likelihood ratio test of θ: 2.22 (0.07)	
Model statistics			
Log likelihood	−2831.94		
Wald χ^2 (Prob > χ^2)	91.28 (0.00)		
Number of failures	543		
Number of observations	543		
Number of groups (states)	28		
Observations per group	Min: 2; Max: 72; Avg.: 19.4		
Global PH test:	χ^2: 9.6 (Prob > χ^2): 0.65		

Source: The 2004 American National Election Study.
***$p < 0.01$, **$p < 0.05$, *$p < 0.1$, *two-tailed*.
[a]We use the Efron method for ties. The multilevel model is estimated using a Gamma distribution for the frailty.

[11] We ran a second version of the model without a frailty term, and the estimates for the parameters and their standard errors are comparable across the models (though there are small differences on about ½ of the covariates (results not shown)).

though the likelihood ratio test shows that the random effect is statistically significant. Thus, there is significant within-state correlation/"state-wise" heterogeneity.[12]

Figures 32.1 and 32.2 present the "backed out" cumulative hazard and baseline hazard rates for the Cox model with a shared frailty; although the Cox model does not parameterize the baseline, it can be retrieved after estimation. Looking at the figures, we see that slope in the cumulative hazard function changes with time, and that the baseline hazard has a steeper slope near the end of the election cycle, but lowers for roughly days 225 – 250. The results show that voters have a higher risk of deciding starting approximately 135 days into the campaign, and that this risk continues to increase until it peaks about a month before election day. In general, these results suggest there is duration dependence in the data. For all observations, events at some time periods are at more risk to occur in comparison to other time periods. Furthermore, this rate appears to fluctuate with no specific functional form.

Lastly, Fig. 32.3 shows the "electoral frailty" of American states in the 2004 presidential contest, as we graph the group-wise frailty estimates for the states included in the sample.[13] From this graph, we can see that states above 0 are the most failure prone (the most frail state was Minnesota, with a value of 0.11), and those below the line are the least failure prone (the least frail state was Texas, home of the incumbent, with a value of −0.26). Our analysis

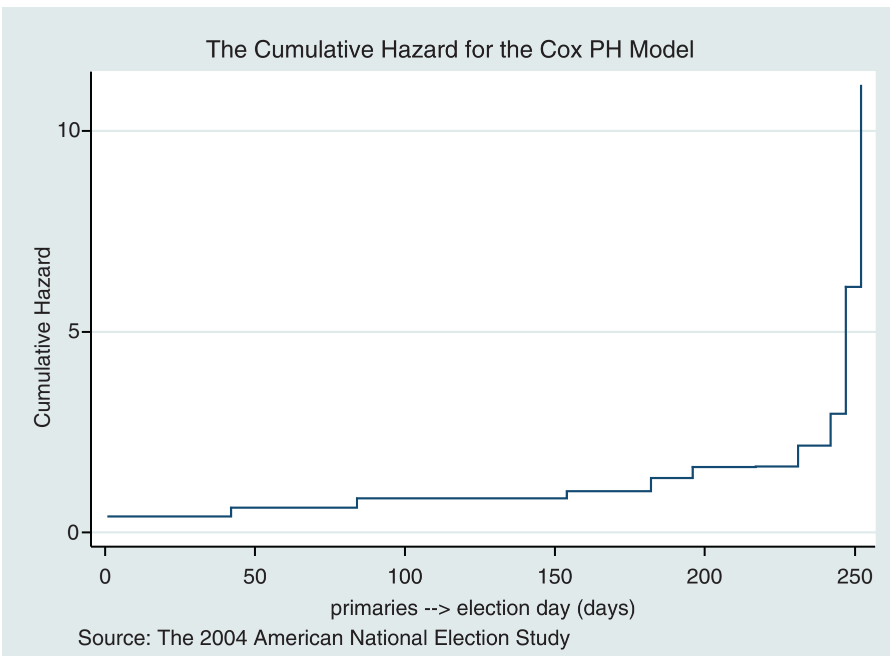

Figure 32.1. The Cumulative Hazard for the Multilevel (Shared Frailty) Cox Model

---

[12] When estimating a shared frailty model, we must be careful to note that the interpretation of the hazard ratios becomes conditional on the frailty term (Cleves et al. 2004). However, as the frailty term – theta – approaches 0, the interpretation of the hazard ratios returns to normal. In the case of our model, we do not concern ourselves with additional interpretation given the relatively small size of the frailty effect.

[13] The x-axis is marked according to ICPSR state identification numbers, and bears no relevance to the analysis other than to enable the graphical presentation of the frailty estimates.

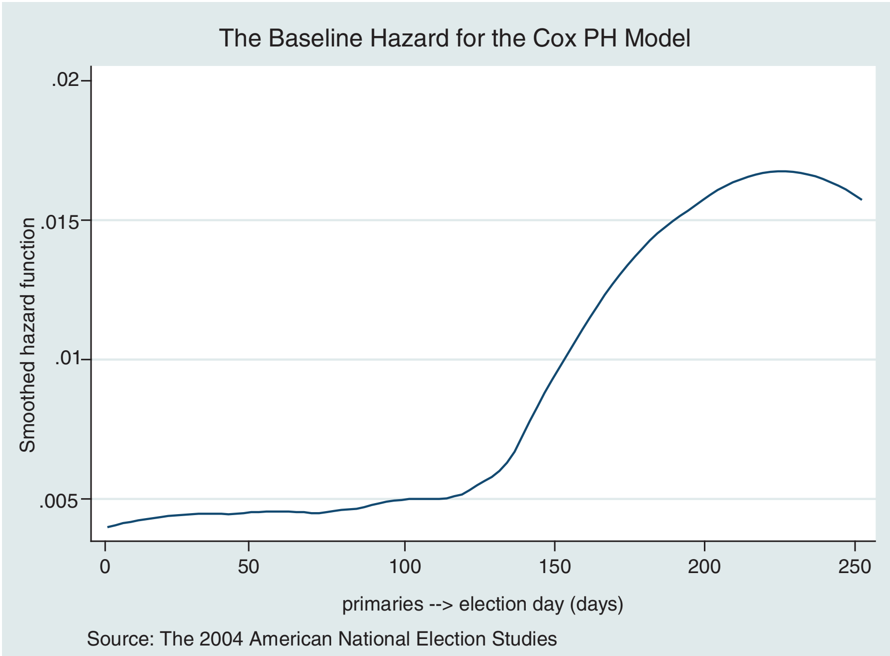

FIGURE 32.2. The Baseline Hazard for the Multilevel (Shared Frailty) Cox Model

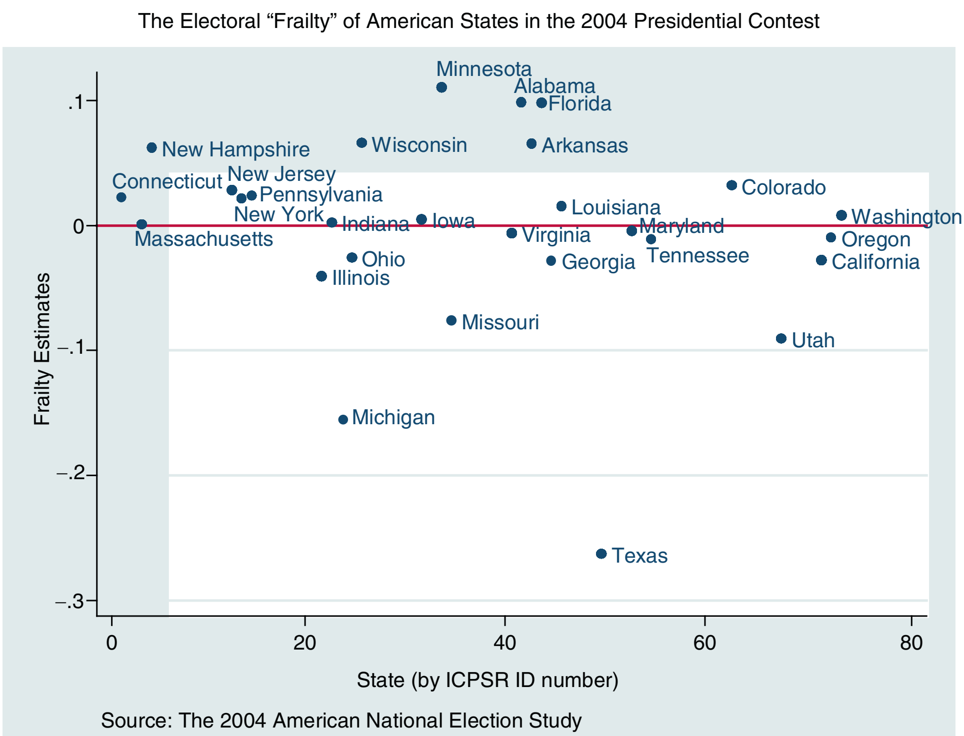

FIGURE 32.3. The Electoral "Frailty" of American States in the 2004 Presidential Contest. Least Frail State: Texas (–0.26). Most Frail State: Minnesota (0.11)

indicates that although Bush carried Texas by a comfortable margin, Texans were more likely to decide later in the election.

DISCUSSION AND CONCLUSION

Event history leverages temporal information when social scientists find that the answer to "why" necessitates an answer to "when," and the result is a more nuanced understanding of process and (ultimately) the subject under study – one that is more empirically and theoretically satisfying. Event history analysis accounts for important temporal dynamics.

There is a growing body of longitudinal data available as social scientists have come to recognize the increased inferential leverage of event history analysis. In particular, within the last decade, the use of these techniques has ballooned in the study of politics due to increased interest in the concepts of survival and risk – processes that are an inherent part of survival modeling. Thus, event history analysis often provides an ideal example of synergy between method and question.

The future of event history in the study of politics is promising, as approaches and extensions continue to provide less restrictive and more realistic accounts of the dynamic, longitudinal processes being studied. As we noted previously, most recently these extensions have included the ability to account for heterogeneity, event dependence, and increasingly spatial relationships, and each of these extensions are exciting and powerful in their own right. Heterogeneity may arise in a number of contexts, perhaps as subjects vary in their ability to learn, leaders take differential risks, and cultures diverge (Box-Steffensmeier et al. 2007)[14]; event dependence arises from the repeated occurrences of events, and has long been recognized as a concern in the study of politics (Andersen and Gill 1982; Box-Steffensmeier and Zorn 2002; Wei et al. 1989; Prentice et al. 1981). And lastly, the incorporation of spatial dependencies (though most recent) is also of great interest to those studying politics (Banerjee et al. 2003; Berry and Baybeck 2005; Volden 2006; Boehmke 2007; Darmofal 2007) – after all, the diffusion of policies across states is likely to impact the timing of the adoption of policies, just as the proximity of civil unrest is likely to impact the timing of outbreaks of additional violence. Researchers are discovering both new questions *and* new answers to old questions when conducting event history analyses, which bodes well for the social sciences in terms of disciplinary progress and maturity.

REFERENCES

Andersen, Per Kragh, and R.D. Gill. 1982. "Cox's Regression Model for Counting Processes: A Large Sample Study." *The Annals of Statistics* 10: 1100–20.

Andersen, P.K., J.P. Klein, and M. Zhang. 1999. "Testing for Centre Effects in Multi-Centre Survival Studies: A Monte Carlo Comparison of Fixed and Random Effects Tests. *Statistics in Medicine*. 18(12)1489–1500.

Banerjee, Sudipto, Melanie M. Wall, and Bradley P. Carlin. 2003. "Frailty Modeling for Spatially Correlated Survival Data, with Application to Infant Mortality in Minnesota." *Biostatistics*. 4(1): 123–42.

Beck, Nathaniel. 1999. "Modelling Space and Time: The Event History Approach." In *Research Strategies in Social Science*, Elinor Scarbrough and Eric Tanenbaum, eds., Oxford University Press.

[14] See also Sastry (1996) and Kuate-Defo (2001), who explicitly incorporate heterogeneity into their event history models of childhood mortality.

Beck, Nathaniel, Jonathan Katz, and Richard Tucker. 1998. "Taking Time Seriously: Time-Series-Cross-Section Analysis with a Binary Dependent Variable." *American Journal of Political Science*, Vol. 42, 1260–1288.

Berry, William, and Brady Baybeck. 2005. "Using Geographic Information Systems to Study Interstate Competition." *American Political Science Review*. 99: 505–19.

Blossfeld, Hans-Peter, and Gotz Rohwer. 2002. *Techniques of Event History Modeling*. 2nd ed. Nahwah, NJ: Lawrence Erlbaum Associates.

Boag, J.W. 1949. "Maximum Likelihood Estimates of the Proportion of Patients Cured by Cancer Therapy." *Journal of the Royal Statistical Society*, *Series B* 11: 15–44.

Boehmke, Fred. 2007. "Potential Ambiguities in a Directed Dyad Approach to State Policy Formation." Presented at the Annual Meeting of the Political Methodology Society. State College, PA.

Box-Steffensmeier, Janet M., and Christopher J.W. Zorn. 2002. "Duration Models for Repeated Events." The Journal of Politics. 64 (November):1069–94.

Box-Steffensmeier, Janet M., and Suzanna De Boef. 2006. "Repeated Events Survival Models: The Conditional Frailty Model." *Statistics in Medicine*. 25(20, October): 3518–3533.

Box-Steffensmeier, Janet M., Suzanna De Boef, and Kyle Joyce. 2007. "Event Dependence and Heterogeneity in Duration Models: The Conditional Frailty Model." *Political Analysis*.

Box-Steffensmeier, Janet M., Suzanna De Boef, Banks Miller, and Anand Sokhey. 2008. "Repeated Events Data and the Conditional Frailty Model: A Foster Care Application." Ohio State University, Working Paper.

Box-Steffensmeier, Janet and Brad Jones. 1997. "Time is of the Essence: Event History Models in Political Science." *American Journal of Political Science*. 41: 1414–61.

Box-Steffensmeier, Janet and Brad Jones. 2004. *Event History Modeling: A Guide for Social Scientists*. New York: Cambridge University Press.

Box-Steffensmeier, Janet M., Peter Radcliffe, and Brandon Bartels. 2005. "The Incidence and Timing of PAC Contributions to Incumbent U.S. House Members, 1993–94." *Legislative Studies Quarterly*. 30 (4, November): 549–79.

Box-Steffensmeier, Janet M. and Christopher J.W. Zorn. 2001. "Duration Models and Proportional Hazards in Political Science. *The Journal of Politics*. 64: 1069–94.

Clark, David and Patrick Regan. 2003. "Opportunities to Fight." *Journal of Conflict Resolution*. 47(1): 94–115.

Cleves, Mario A., William W. Gould, and Roberto G. Gutierrez. 2004. *An Introduction to Survival Analysis Using Stata, revised edition*. College Station, TX: Stata Press.

Collett, D. 1994. *Modeling Survival Data in Medical Research*. London: Chapman and Hall.

Conover, Pamela and Stanley Feldman. 1981. "The Origins and Meaning of Liberal and Conservative Self-Identifications." *Amer. Journal of Political Science*. 25: 617–645.

Cox, D.R. 1972. "Regression Models and Life Tables." *Journal of the Royal Statistical Society. B* 34: 187–220.

Cox, D.R. 1975. "Partial Likelihood." *Biometrika* 62: 269–76.

Darmofal, David. 2007. "Bayesian Spatial Survival Models for Political Event Processes." Typescript. University of South Carolina.

Delli Carpini, Michael X. and Scott Keeter. 1996. *What Americans Know About Politics and Why It Matters*. New Haven, CT: Yale University Press.

Diermeier, Daniel, and Randy T. Stevenson. 1999. "Cabinet Survival and Competing Risks." *American Journal of Political Science*. 43: 1051–68.

Downs, A. 1957. *An economic theory of democracy*. New York: Harper and Row.

Fiorina, M.P. 1981. *Retrospective voting in American national elections*. New Haven: Yale University Press.

Garibotti, Gilda, Ken R. Smith, Richard A. Kerber, and Kenneth M. Boucher. 2006. "Longevity and Correlated Frailty in Multigenerational Families." *Journal of Gerontology*. Vol. 61A, No. 12: 1253–61.

Golub, Jonathan, and David Collett. 2002. "Improving Model Specification in Duration Analysis: Time-sensitive Covariates and Cox Models." Working paper, Reading University.

Grambsch, P.M. and T.M. Therneau. 1994. "Proportional Hazards Tests and Diagnostics Based on Weighted Residuals." *Biometrika* 81:515–526.

Harrell, F.E. 1986. "The PHGLM Procedure." *SUGI Supplemental Library User's Guide*. Cary, NC: SAS Institute.

Huckfeldt, Robert, and John Sprague. 1995. *Citizens, Politics, and Social Communications: Information and Influence in an Election Campaign*. New York: Cambridge University Press.

Jones, Daniel M., Stuart A. Bremer, and J. David Singer. 1996. "Militarized Interstate Disputes, 1816–1992: Rationale, Coding Rules, and Empirical Patterns." *Conflict Management and Peace Science*. 15: 163–213.

Key, V.O., Jr. 1966. *The responsible electorate*. Harvard University Press.

King, Gary, James E. Alt, Nancy E. Burns, and Michael Laver. 1990. "A Unified Model of Cabinet Dissolution in Parliamentary Democracies." *American Journal of Political Science*. 34: 846–871.

King, Gary and Langche Zeng. 2001. "Explaining Rare Events in International Relations." *International Organization*. 55: 693–715.

Kuate-Defo, Barthelemy. 2001. "Modelling Hierarchically Clustered Longitudinal Survival Processes with Applications to Childhood Mortality and Maternal Health." *Canadian Studies in Population*. 28(2): 535–61.

Li, Y., and X. Lin. 2003. "Functional Inference in Frailty measurement Error Models for Clustered Survival Data Using the SIMEX Approach." *Journal of the American Statistical Association*. 98(461): 191–204.

Luskin, Robert C. 1987. "Measuring Political Sophistication." American Journal of Political Science 31:856–99.

McClurg, Scott D. 2006. "The Electoral Relevance of Political Talk: Examining Disagreement and Expertise Effects in Social Networks on Political Participation." *American Journal of Political Science* 50(3): 737–754.

Prentice, R.L., B.J. Williams, and A.V. Peterson. 1981. "On the Regression Analysis of Multivariate Failure Time Data." *Biometrika*. 68: 373–9.

Sastry, Narayan. 1996. "Community Characteristics, Individual and Household Attributes, and Child Survival in Brazil." *Demography*. 33(2): 211–29.

Schmidt, Peter and Ann D. Witte. 1988. *Predicting Recidivism Using Survival Models*. New York: Springer-Verlag.

Schrodt, Philip A. 2000. "Pattern Recognition of International Crises using Hidden Markov Models." In *Political Complexity: Nonlinear Models of Politics*, ed. Diana Richards. Pp. 296–328. Ann Arbor: University of Michigan Press.

Therneau, Terry M., and Patricia M. Grambusch. 2000. *Modeling Survival Data*. New York: Springer.

University of Michigan, Center for Political Studies. *AMERICAN NATIONAL ELECTION STUDY, 2004: PRE- AND POST-ELECTION SURVEY* [Computer file]. ICPSR04245-v1. Ann Arbor, MI: University of Michigan, Center for Political Studies, American National Election Study [producer], 2004. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2006–02–17.

Volden, Craig. 2006. "States as Policy Laboratories: Emulating Success in the Children's Health Insurance Program." *American Journal of Political Science*. 50: 294–312.

Warwick, Paul. 1992. "Economic Trends and Government Survival in West European Parliamentary Democracies." *American Political Science Review*. 86: 875–86.

Wei, L.J., D.Y. Lin, and L. Weissfeld. 1989. "Regression Analysis of Multivariate Incomplete Failure Time Data by Modeling Marginal Distributions." *Journal of the American Statistical Association*. 84: 1065–73.

Werner, Suzanne. 2000. "The Effects of Political Similarity on the Onset of Militarized Disputes, 1816–1985." *Political Research Quarterly*. 53(2): 343–74.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.