Congrès AFSP Toulouse 2007

Table ronde 1
"Réflexions sur les méthodes en science politique des deux côtés de l'Atlantique"

Session 2
La prise en compte du temps

EVENT HISTORY ANALYSIS: A LONGITUDINAL APPROACH FOR VOTING DECISIONS

Janet M. Box-Steffensmeier
Vernal Riffe Professor of Political Science
Director, Program in Statistics & Methods
The Ohio State University
2140 Derby Hall, 154 N. Oval Mall
Columbus, OH 43210-1373
Steffensmeier.2@osu.edu

Anand E. Sokhey
Fellow, Program in Statistics & Methods
Ph.D. Student, Dept. of Political Science
The Ohio State University
2140 Derby Hall, 154 N. Oval Mall
Columbus, OH 43210-1373
Sokhey.2@osu.edu

ABSTRACT

In this manuscript we focus on event history analysis, noting several prominent applications to the study of politics. We begin by discussing different modeling strategies, along with problems and misconceptions common (and unique) to political survival research. We then introduce the Cox proportional hazards model, describing its logic, estimation, and differences when compared to parametric approaches.

In the second portion of the paper, we highlight important extensions to the Cox approach, focusing on multi-level (hierarchical) and conditional frailty event history models. Using the 2004 American National Election Study, we present a basic example of a multi-level Cox model (in which individual voters are nested in states) that examines the timing of voting decisions – that is, we ask "When in a campaign do voters make up their minds?" We close with a brief commentary on the future of event history analysis in the study of politics.

Paper prepared for the 2007 meeting of the French Political Science Association, Toulouse, France, September 5-7. Authors are listed alphabetically. Name of the electronic file: TR1sess1Name.doc

An Introduction to Event History Analysis

Researchers are often interested in more than just the occurrence or non-occurrence of a political events; frequently the *timing* of said events is of equal substantive importance, whether it's the dissolution of a government's cabinet (King, Alt, Burns, and Laver 1990), the presence of international military disputes (King and Zeng 2001), contributions by political action committees (Box-Steffensmeier, Radcliffe, and Bartels 2005), or as we will examine in this paper, when a voter makes up her mind in an election campaign. Examining when an event occurs provides additional information and may lead to additional insight about the event. Event history – or survival analysis – is the tool of choice when political scientists find that the answer to "why" necessitates an answer to "when."

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 her mind before she decides who 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." Researchers may deal with data processes in which there either are multiple failures (i.e., repeated events), multiple spells (i.e., periods during which a subject is at risk of failing), or both multiple failures and spells. These additional data complications are straightforward to address.

As 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:

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

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

$$f(t) = dF(t)/d(t) \qquad \textbf{(2)}$$

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

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

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*, $h(t)$,

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

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. Political 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.

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 the factors of 1) censoring and 2) time varying covariates; OLS is an improper technique for modeling failure-time processes due to its inability to deal with these issues. Censoring occurs when an observation's full history is not observed. For example, in studying the duration of an international military dispute, the dispute may be ongoing at the end-time of the analysis, re: it has not ended, and thus the dispute is *right censored*. Event history analysis is specifically designed to account for censored data via the calculation of the hazard rate. *Left-truncation* occurs when some observations have experienced an event before the beginning of the study; it can also be considered a censoring problem as data are not observed, only in this case the non-observation occurs prior to the start of the study. Time-varying covariates are also readily incorporated into the study of event history data, and 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 assess whether war chests impact whether a challenger enters an electoral race –war chests need to be measured over the course of the election cycle as simply measuring war chests 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: chapter 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 – a relative of the exponential – might not be a bad decision.[1] Other parametric models such as the log-logistic can offer the researcher a bit more flexibility in that they allow the specification of non-monotonic hazard rates.

---

[1] 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.

All such parametric models are estimated through maximum likelihood, with the likelihood function having been 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 non-monotonic (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, a test 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. Both parametric models and the Cox model are perfectly acceptable ways to proceed with event history estimation. However, the Cox model does offer some advantages. The Cox approach is a more straightforward and a powerful alternative to parameterization techniques, for 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).[2]

In the Cox model, the hazard rate for the *i*th individual is specified as:

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

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

[2] Discrete and continuous time approaches are also both perfectly acceptable ways to proceed with event history estimation. However, the continuous time approach – which we discuss here – is again more straightforward. By using a continuous time approach, one does not have to fit a link function for the duration dependence. See, however, Beck, Katz, and Tucker (1998) and Beck (1999) who argue that discrete time approaches are more straightforward to interpret due to researchers' familiarity with discrete time (re: logit and probit) models.

"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 in the example presented below.

Unlike the aforementioned models, 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 (something that makes sense, as the Cox model does not directly model the duration dependence in the data). Rather, it is the ordered failure times that contribute information to the partial likelihood function – time only matters to the extent that it gives order to the failure times (Cleves et al. 2004; Collett 1994).

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 *j*th 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).

$$\Pr(t_j = T_i \mid R(t_i)) = \frac{e^{\beta' x_i}}{\sum_{j \in R(t_i)} e^{\beta' x_j}} \quad \textbf{(6)}$$

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

$$\mathcal{L}_P = \prod_{i=1}^{K} \left[\frac{e^{\beta' x_i}}{\sum_{j \in R(t_i)} e^{\beta' x_j}} \right]^{\delta_i} \qquad \textbf{(7)}$$

Given that the Cox model's partial likelihood function is based solely on the ordered failures in the data, estimation cannot take place in the presence of "ties" – or coterminous events – unless the risk set is approximated through other means. The issue of ties is relevant for any continuous time model. Fortunately, researchers have developed several methods for dealing with this problem. The Cox approximations – such as the Breslow, Efron, and Exact Discrete – have greatly improved as a result of the increase in computing power. Indeed, another advantage of the Cox model over other 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 models (multi-level 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; 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 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, De Boef, Miller, and Sokhey 2007).

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. Therneau and Grambusch (2000) define a frailty (or random effect) as a continuous variable that describes excess risk (or frailty) 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 district, state, country, or region. It is also worth noting that levels may be defined by distance rather than 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 overlap (see Banerjee, Wall, and Carlin 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, observations) within a group. The multi-level Cox model is another useful extension that should be of great interest to social scientists.

The conditional frailty model is a Cox model extension designed to account for the presence of both repeated events and heterogeneity. The basic Cox model assumes that the baseline hazard is constant across all observations, and imposes the restriction that the dependent variable is constant conditional on the included covariates. In contrast, with the conditional frailty stratification by event number provides the flexibility of varying baseline hazards (by event number) 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 is more realistic for most datasets (see Box-Steffensmeier and DeBoef 2006; Box-Steffensmeier, DeBoef, and Joyce 2007).[3]

WHEN DO VOTERS MAKE UP THEIR MINDS?

THE DATA: THE 2004 AMERICAN NATIONAL ELECTION STUDY

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.[4] The specific wording of the 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 asks who the respondent voted for in the presidential election.

The general categories and distribution of responses for this question are provided in Table 1. Looking at the table, we note that 33 percent of respondents stated that they knew "all along" how they would vote. At the other end of the spectrum, over 15 percent reported deciding within the last two 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.

[Insert Table 1 About Here]

To the best of our knowledge, no one has looked at this question – i.e., the timing of the voting decision – in quite this way.[5] 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 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.[6] 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 on 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 conversation (with family,

[3] Future work will explore correlated, rather than shared, frailty such that the random error of unobserved heterogeneity is common among multiple occurrences of the event for each individual, country, etc.
[4] The 2004 American National Election Study is available through the Inter-University Consortium for Political and Social Research (ICPSR).
[5] 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.
[6] 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).

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 experience leads a voter to an earlier decision on whom to support in the election. Experienced voters may reach a decision sooner – something which may reflect 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 on 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).

Finally, we include a number of demographic controls to test whether there are differences by age, gender, education, income, and race. We do not expect such differences to emerge, though, as there is no theory 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.

RESULTS

Researchers should always conduct diagnostic tests to check for problems, and thus we begin by examining a few common procedures as they pertain to our model: the link test, and the 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, a link test does not reveal any concerns as the p-value is not statistically significant. Specifically, 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.

[Insert Table 2 About Here]

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 model, 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. The graphical tests *are* still useful, however, particularly when trying to determine which function of time to use as an interaction with offending covariates when fixing any violations of the proportional hazards assumption. Our tests here show that there is not a problem with the proportional hazards assumption.

Looking at the results (please see Table 2), we find that both 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. For hazard ratios, 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 table, we see that the estimated size of the random effect (akin to a random effect in a multi-level model) is 0.021, which is small – though the likelihood ratio test shows that the random effect is statistically significant. This gives us confidence that there is significant within-state correlation/"state-wise" heterogeneity.[7]

Figures 1 and 2 present the "backed out" cumulative hazard and baseline hazard rates for the Cox model – indeed, 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 to 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 around 225 days into the campaign. 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, figure 3 shows the electoral frailty of American states in the 2004 presidential contest. That is, below we graph the group-wise frailty estimates for the states included in the sample (the numbers next to the state names are simply survey ID

[7] 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 worry much about additional interpretation given the relatively small size of the frailty effect.

numbers, and bear no relevance to the analysis). 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 indicates that although Bush carried Texas by a comfortable margin, Texans were more likely to decide later in the election.

[Insert Figures 1-3 About Here]

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 – of the subject under study that is more empirically and theoretically satisfying. In addition, by taking into account temporal dynamics, survival analysis can help prevent scholars from jumping to misleading conclusions about relationships under consideration.

The future of event history in the study of politics is promising, as modeling extensions continue to provide less restrictive and more realistic accounts of the dynamic, longitudinal processes being studies – most recently these extensions have included the ability to account for heterogeneity, event dependence, and spatial relationships. In addition, there is better (and more) data now available for use with event history techniques; in turn, better understandings of event history continue to lead directly to the collection and creation of suitable survival data sets. Indeed, social scientists are discovering new answers to old questions when conducting (what is often *more* appropriate) analyses using event history approaches.

Table 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/9months 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)/3-4months"	50	6.09	154
"at the time of the Republican convention (8/30-9/2/04)/2-3months"	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
"~two 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 proceeding asks who the respondent voted for in the presidential election. Source: The 2004 American National Election Studies			

Table 2:
Predicting the Timing of Voting Decisions, 2004
(Cox Proportional Hazards Estimates; Multi-Level (Shared Frailty) Model)†

Covariates	*Coefficient*	*Hazard Ratio*	*S.E.*	
Personal Political Characteristics				
Strength of ideology	.21	1.24	.05	***
Strength of partisanship	.24	1.27	.05	***
Disapprove of the President	-.01	.99	.09	
Political Engagement				
Frequency of Political Discussion	.02	1.02	.02	
Voted in 2000	.24	1.27	.14	*
Political Interest	.35	1.42	.09	***
Political Knowledge	.07	1.07	.05	
Demographics				
Age	-.00	1.00	.00	
Gender	.04	1.05	.09	
Education	-.01	.99	.02	
Income	.00	1.00	.01	
Race (African American)	.24	1.28	.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*				
†Efron method for ties. Gamma frailty.				

Figures 1 and 2: The Cumulative Hazard and Baseline Hazard for the Multi-Level (Shared Frailty) Cox Model

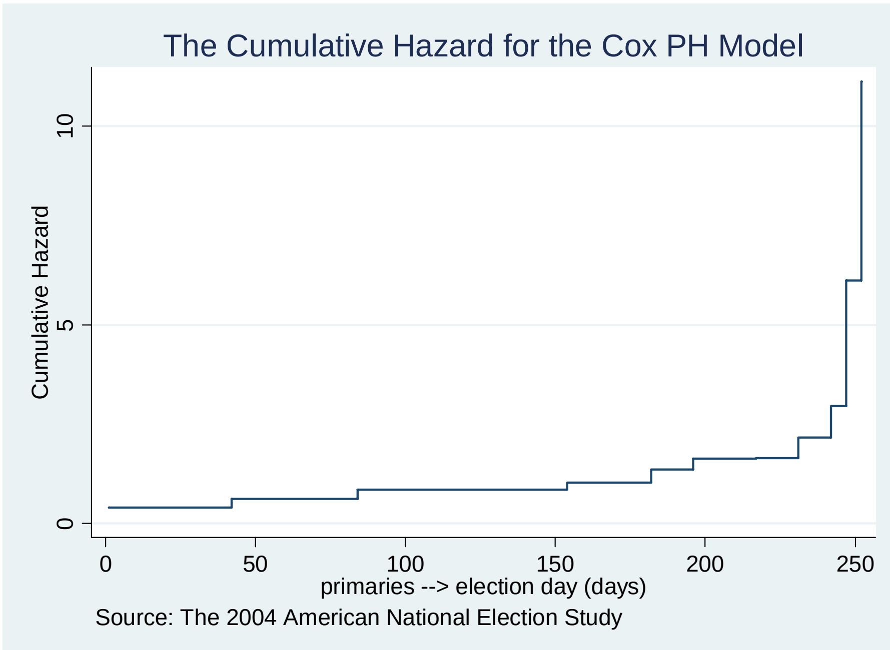

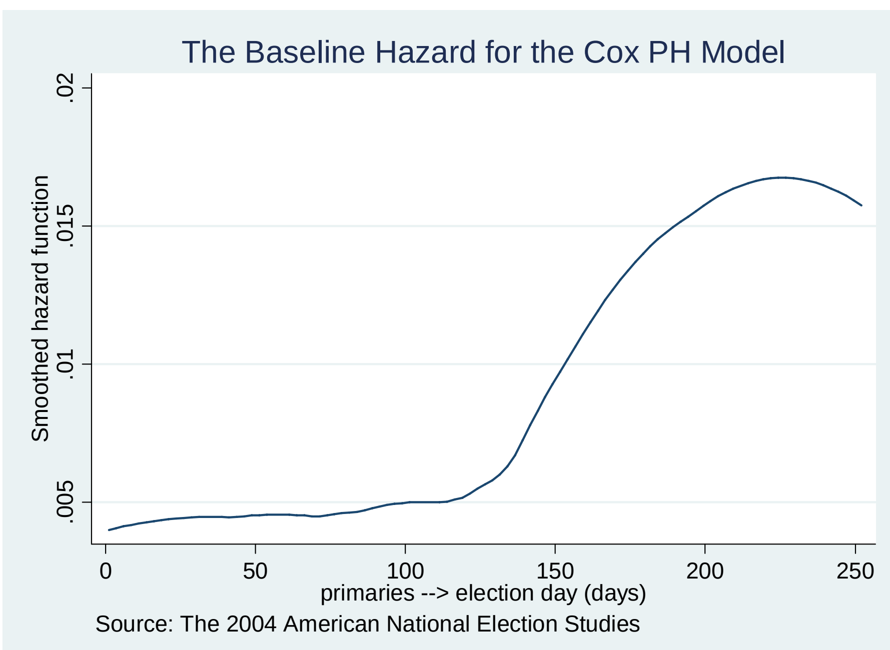

Figure 3: The Electoral "Frailty" of American States in the 2004 Presidential Contest

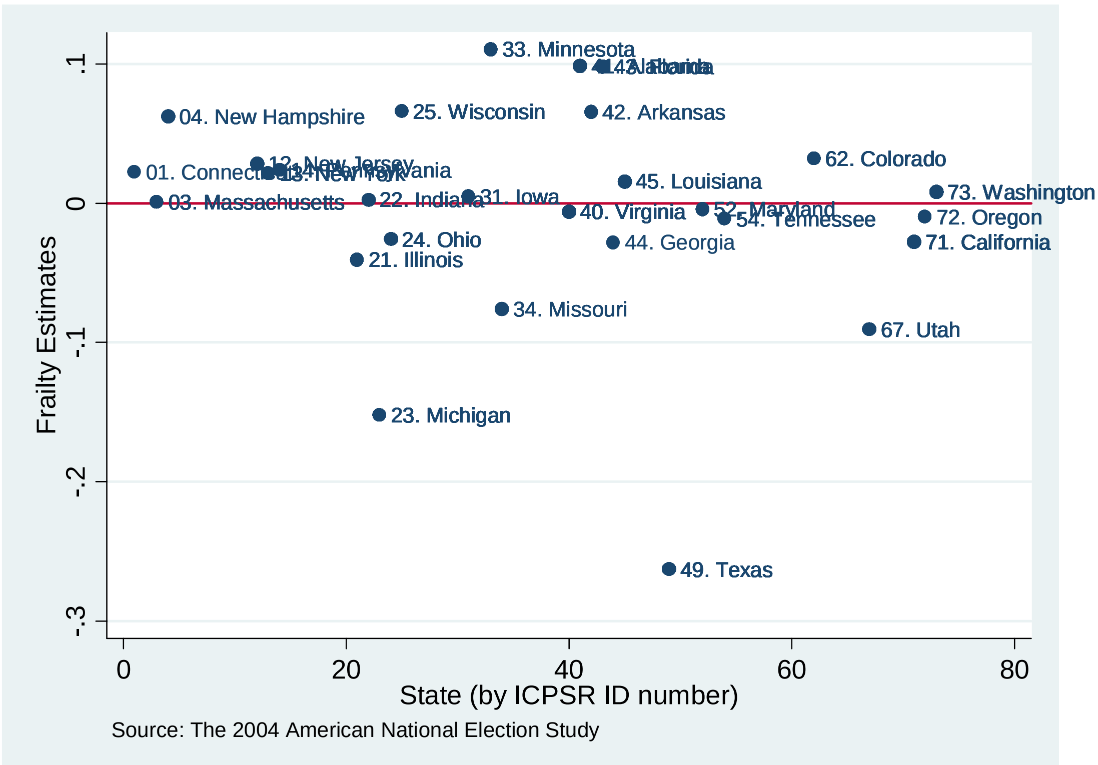

Least Frail State: Texas (-.26)
Most Frail State: Minnesota (.11)

REFERENCES

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.

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.

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

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, Jant M., Suzanna De Boef, Banks Miller, and Anand Sokhey. 2007. "Repeated Events Data and the Conditional Frailty Model: A Foster Care Application." Ohio State University, Working Paper.

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.

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.

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

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

---. 1975. "Partial Likelihood." *Biometrika* 62: 269-76.

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.

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.

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

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.

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