Repeated Events

1 Introduction

To date, we have either (i) assumed that when an observation has an event (or fails) that it leaves the data set or (ii) ignored the fact that an observation can experience more than one event.[1] However, in many political science settings, it is possible for the same unit to experience repeatable events. For example, dyads can experience multiple wars. The issue of repeated events is yet again an issue with dependence: second and subsequent events are likely to be influenced by, and therefore different from, first events. Political actors are likely to learn from past events and past events can put actors done a path where options can become more or less constrained. Such dependence is called ‘subject event dependence’. A result of all of this is that treating ‘repeated events’ as though they were independent events runs the risk of yielding misleading results for two reasons. First, the presence of correlated events presents a problem similar to autocorrelation in conventional regression analysis: by treating each observation as independent, we overstate the information that we have, leading to incorrect estimates of standard errors. Second, such models implicitly restrict the influence of covariates to be the same across events, when, in fact, there may be varying effects from one event to the next. There are are a variety of ways to deal with repeated events (Box-Steffensmeier & Zorn 2002, Cleves 1999).

The approaches to dealing with repeated events are summarized in Figure 1. All of these approaches

Figure 1: Models for Repeated Events

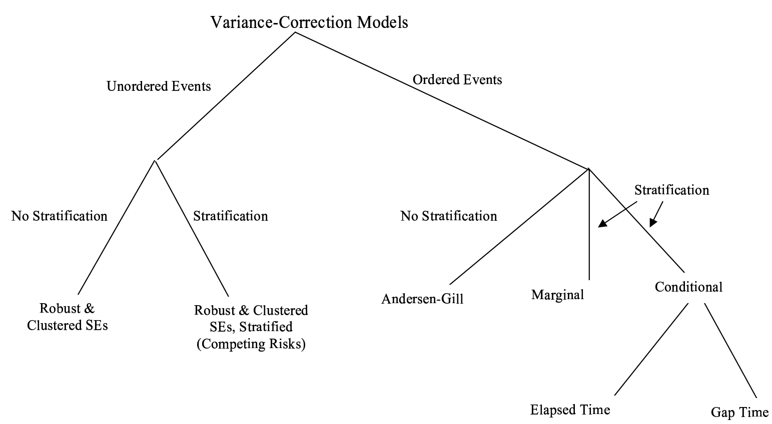

fall under the heading of ‘variance-correction models’.[2] The idea is that repeated events only affect the variance of estimates, not the means. If there’s dependence, then the $\hat{\beta}$s are still consistent (but biased) - it’s the estimate of the variance-covariance matrix that is biased and inconsistent. So, the idea is that you

[1] When we looked at discrete time models of dispute, Beck, Katz and Tucker (1998) did put in a counter for the number of previous disputes. This was certainly better than nothing and the simplest thing that you can do. The event counter allows the baseline hazard to increase by a set proportion with each subsequent dispute. But, note that a fixed monotonic change across events is a strong assumption and may not really capture the effects of repeated events. We are now going to look at alternative ways to deal with repeated events.

[2] It is also possible to deal with repeated events by taking a frailty model approach. The idea is that there is possible dependence across events within a subject that is captured by a frailty term ν_i. We’ve already seen how to estimate these models. This approach makes some sense if you think that the frailty is fixed over time.

simply estimate the model treating events as independent and then 'fix' the variance estimates after the fact. This is analogous to OLS methods for dealing with heteroskedasticity/autocorrelation. Essentially, variance-correction models differ along three dimensions:

1. **The Risk Set – Are events ordered or unordered?**

   Variance-correction models differ in how they define the risk set. As before, the risk set defines how subjects are considered to be at risk of experiencing a given event k at a particular time t. Obviously, when defining the risk set, it is important to know whether the repeatable events follow a sequential order or whether they can occur simultaneously or out of order. If a subject cannot experience a subsequent event without having experienced a prior event, then we need to define the conditional risk set that preserves this ordering. If there is no ordering, then the risk set can be 'unconditional' and we can estimate a 'marginal' model.

2. **Clock Time**

   Variance-correction models differ in how they count time at risk for an event. Essentially the difference is between whether (i) the clock begins when a unit enters the observation period and ends when a subject is censored (elapsed time) or (ii) whether the clock restarts after each event (gap time). In elapsed time models, a unit experiencing events in months 2, 8, and 12 would have start-stop times of 0-2, 0-8, and 0-12. In gap time models, the same unit would have start-stop times of 0-2, 0-6, and 0-4.

3. **Stratification**

   Variance-correction models differ in terms of whether they allow for event-specific baseline hazards. If the hazard rate is likely to vary across events – as when event dependence exists – then we would want a model that is able to stratify the data to allow for the estimation of a separate baseline hazard for each event.[3] If the risk associated with all events is the same, then a single baseline hazard will work and stratification is not necessary.

Basically, you have to think about the data generating process for your particular situation and determine where you fall on each of these three dimensions. This will then help you determine which particular repeated events model is suitable for your purposes. Depending on the risk set, the clock time, and stratification, you can estimate very different dynamic processes. Box-Steffensmeier, De Boef, and Joyce (2007) argue that a conditional-risk set model estimated in gap time with stratification is the most natural variance-correction model for repeated events in political science.

2 Unordered Events

Unordered events are events that can occur to an observation but where the sequence does not matter. It is actually quite hard to think of examples of this in the political science setting.[4]

[3] As we saw with the stratified Cox model for competing risks, each event can have a different baseline hazard but we assume that the effects of the covariates are constant across the different strata. To get around this, we might interact the covariates with the strata variables.

[4] For more detail on how to estimate all of these models in STATA, see Cleves (1999).

2.1 Same Type – No Stratification

If you think that the unordered events are the same, then you can deal with this by using robust standard errors clustered on the unit of analysis - this deals with the temporal dependence in the data.

```
stset duration, failure();

stcox  X, cluster() efron nohr;
```

Note that you don't actually need to specify robust standard errors since STATA knows to produce robust errors when the cluster command is used.

2.2 Different Type – Stratification

If the unordered events are different, then we are really in the world of competing risks where events can be repeated. In other words, a patient might be at risk of suffering from different illnesses in no systematic order. To take this situation into account, we would stratify the observations on failure type, then allow each stratum to have its own baseline hazard function but restrict the coefficients to be the same across the strata. An important characteristic of these failure events is that each event can occur only once per subject. All subjects are at risk for all events at the beginning. If a subject suffers one event, then he remains in the data set but is only at risk for $n-1$ further events. This is basically the stratified Cox model that we looked at to deal with competing risks.

```
stset duration, failure();

stcox  X, cluster() strata() efron nohr;
```

3 Ordered Events

The principal difference between the ordered event models has to do with the way in which the risk sets are defined at each failure time.

3.1 Andersen-Gill

For each unit, there must be one data record per event or time interval. This means that if a unit experiences one event, then it will have two observations - the first covers the time span from entry in the data set to the time of the event and the second observation spans the time from the event to the end of the data set and is censored.

```
stset duration, failure() exit( .) id(id) enter();

stcox  X, robust efron nohr;
```

Again, you will need to use robust clustered standard errors. However, because you need to specify ID() when you STSET the data, asking STATA to get robust standard errors by typing ROBUST automatically gets STATA to cluster them on the units specified in ID(). The primary characteristic of the AG model is the assumption that the risk of an event for a given unit is unaffected by any earlier events that occurred to the same unit unless covariates that capture such dependence are explicitly included in the model. Thus, the AG model specifies an unconditional risk set in elapsed time without stratification. As you should have noticed, the AG model and the Cox model are identical except that the AG model uses robust standard errors clustered within units. That the AG model assumes that the underlying baseline hazard is the same for all events helps to distinguish it from the other 'ordered' models we will look at. The AG assumption that the event times are conditionally independent is quite a strong assumption and hard to justify in most repeated event settings.

3.2 Marginal Risk Set Model

In this approach, we employ the traditional competing risks setup for multiple events to repeated events. Ordered events data are treated as if they presented a typical competing risks problem: each observation is at risk for the first, second, third etc. even from the beginning of the study period. The data are then stratified by event number and separate baseline hazards are estimated at each occurrence. The signature characteristic of this approach is that all observations are at risk for all events at all times prior to experiencing that event. This means that, say, the fifth event can (in theory) occur at any time, even prior to the 'first', 'second' etc. events. In other words, there is no natural sequence to the events here. Whether this is plausible will depend on your data generating process.

```
stset duration, failure();

stcox  X,  efron nohr strata() cluster();
```

Again, you are using robust standard errors clustered on the different strata. This approach is different from Andersen-Gill in that you get different baseline hazards for each event occurrence. However, it is the same in that you get a single set of covariates that are constant across the event ranks. You could include interactions between your covariates and strata to allow the covariates to vary by strata. Essentially, the marginal risk set model is an unconditional risk set model in elapsed time with stratification.

Tha AG model and marginal risk set model assume that events are ordered but don't actually model this in a direct way. The only difference between these models and the unordered models that we looked at earlier has to do with how we STSET the data.

3.3 Conditional Risk Set Model

In contrast to the previous models, the conditional risk set model specifically assumes that an observation is not at risk for a later event until all prior events have already occurred. Thus, the risk set at time t for the k^{th} occurrence of an event is limited to those observations under study at time t who have already experienced $k-1$ events of that type. Estimates are stratified by event rank so that we can get different baseline hazards across event ranks. As in the previous models, though, covariate effects are assumed to be constant across strata. You could get around this as before by including interactions between your covariates and strata to allow the covariates to vary by strata.

The conditional risk set model can be estimated in either (i) elapsed time, where estimates are provided for the effect of covariates on the hazard of the k^{th} event since the beginning of the observation period or (ii) gap time, where estimates are provided for the effect of covariates on the hazard of the k^{th} event since the occurence of the $(K-1)^{th}$ event. To estimate the elapsed time model, type

```
stset duration, fail() exit() id() enter();

stcox X, nohr efron robust strata();
```

To estimate the gap time model, type:

```
stset duration, fail() exit()  enter() exit();

stcox X, nohr efron robust strata() cluster();
```

3.4 Summary

A summary of the various variance-correction models and their assumptions is shown in Table 1.

Table 1: Comparison of Variance-Correction Models for Repeated Events

Model Property	Andersen-Gill	Marginal	Conditional Elapsed Time	Conditional Gap Time
Risk Set for Event k at time t	Independent Events	All subjects that haven't experienced event k at time t	All subjects that have experienced event $k-1$ and haven't experienced event k at time t	
Time Scale	Duration since Entry	Duration since Entry	Duration since Entry	Duration since Previous Event
Robust standard errors (clustered)?	Yes	Yes	Yes	Yes
Stratification by Event	No	Yes	Yes	Yes

4 Results

Let's look at some results from the variance-correction models. We'll use the data on state adoption of anti-obscenity legislation again from Box-Steffensmeier and Jones. There is more than one type of anti-obscenity legislation that states can adopt - thus, we have repeatable events. What might we do with this sort of data? One option would be to simply drop all the second events and model the time to the first event.

```
stcox fundam south collc citideo legislativeideology womleg
        murderrate if sequence==1, nohr efron cluster(statnam);
```

Note that we are still trying to deal with temporal dependence by using robust standard errors clustered by the unit of analysis. The results from this analysis are shown in column 1 of Table 2. Now, let's include all the events, use robust standard errors clustered on the unit to take account of time dependence, but not take account of their repeated nature i.e. treat events as though they were independent. This is basically the Andersen-Gill approach.

```
stcox fundam south collc citideo legislativeideology womleg
        murderrate, nohr efron cluster(statnam);
```

These results are shown in column 2. As you can see, the sample size and the coefficients change a lot. Finally, let's use the conditional risk set model where time is measured in gap time i.e. time from the previous event. The data are already STSET to do this.

```
stcox fundam south collc citideo legislativeideology womleg
        murderrate, nohr efron cluster(statnam) strata(sequence);
```

These results are shown in column 3. You could interpret the coefficients from this model exactly as you would from a normal Cox model. The heterogeneity associated with each different event is dealt with through stratification i.e. we allow each event rank to have its own baseline hazard.

Table 2: Repeated Events

	Single Event Model	Non-Stratified Event Model (A-G)	Conditional Gap-Time Model
Fundamentalists(%)	-0.02	0.18	0.22**
	(0.21)	(0.13)	(0.10)
South	0.79	0.19	0.04
	(0.56)	(0.46)	(0.42)
College Graduates	0.11	0.12**	0.12**
	(0.07)	(0.06)	(0.05)
Legislative Ideology	-0.07***	-0.03***	-0.02***
	(0.01)	(0.01)	(0.01)
Citizen Ideology	0.09***	0.05***	0.04***
	(0.02)	(0.02)	(0.01)
Female Legislators (%)	-0.07*	-0.01	0.00
	(0.04)	(0.03)	(0.02)
Murder Rate	0.24***	0.16**	0.18***
	(0.06)	(0.04)	(0.04)
Observations	226	400	400
Log-Likelihood	-147.23	-447.21	-334.18

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$ (two-tailed); Standard errors are given in parentheses

5 Conditional Frailty Model

Last time we looked at frailty models to deal with unobserved heterogeneity.[5] This time we looked at variance-corrected models to deal with 'subject' event dependence that arises in the repeated events setting. It turns out that variance-correction models also deal with correlation in event times that arise due to heterogeneity. Those variance-correction models that allow for stratification also account for correlation in event times due to event dependence. However, because variance-corrected models do not incorporate the heterogeneity directly at the estimation stage – it is always by correcting the variance-covariance matrix after estimation – the estimates remain biased (but still consistent); evidence suggests that the bias is in a downward direction so that covariate effects will be underestimated.

One issue is that neither the variance-correction models nor the frailty models are consistent with repeated event settings in which there may be both unobserved heterogeneity and event dependence. An alternative is the conditional frailty model that basically combines a conditional gap time repeated event models with an unshared or individual frailty model (Box-Steffensmeier, Boef & Joyce 2007). The conditional frailty model controls for event based dependence through event-based stratification and controls for unobserved heterogeneity through the inclusion of a random effect. The model is formulated in gap time and so parameter estimates indicate how the covariates affect the risk of experiencing the k^{th} event for those observations that have experienced the $(k-1)^{th}$ event. The hazard rate for a particular event k occurring to

[5]The notes in this section are based on Box-Steffensmeier, De Boef, and Joyce (2007).

a specific unit i, $h_{ik}(t)$, for the conditional (unshared) frailty (gap time) model can be written as:

$$h_{ik}(t) = h_{0k}(t - t_{k-1})e^{X_{ik}\beta + W_i\psi} \quad (1)$$

where k denotes the event number, h_{0k} is the baseline hazard for event k, $(t - t_{k-1})$ incorporates the gap time structure so that the hazard gives the risk for event k since the $(k-1)^{th}$ event, and W_i is a frailty term from a probability distribution with a mean of 0 and a variance of 1. The partial likelihood for this model is:

$$\mathcal{L} = \prod_{i=1}^{n} \prod_{k=1}^{K} \left(\frac{e^{X_{ik}\beta + W_i\psi}}{\sum_{i=1}^{n} \sum_{k=1}^{K} Y_{ik} e^{X_{ik}\beta + W_i\psi}} \right)^{d_{ik}} \quad (2)$$

where k refers to the event number, d is a censoring variable equal to 1 if event k is observed and 0 if censored, and Y is an at-risk indicator equal to 1 when a unit is at risk for the current event k and 0 otherwise. This model allows for the possibility that both event dependence and unobserved heterogeneity make important contributions to the hazard of a unit's risk for the recurrence of a particular event.

In order to get good estimates of heterogeneity and event dependence, it is important that you have variation across units in the timing of events and in the number of events experienced. Variation on the number of events experienced provides information on a unit's frailty. And variation on the timing of multiple events provides information on event dependence. When the number of cases in a particular stratum is relatively small, you might want to consider collapsing higher strata.

References

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

Box-Steffensmeier, Janet & Christopher Zorn. 2002. "Duration Models for Repeated Events." *Journal of Politics* 64:1069–1094.

Box-Steffensmeier, Janet, Suzanna De Boef & Kyle A. Joyce. 2007. "Event Dependence and Heterogeneity in Duration Models." *Political Analysis* 15:237–256.

Cleves, Mario. 1999. "How Do I Analyze Multiple Failure-Time Data Using STATA?" *STATACorp* .