Survival Analysis

Mark Lunt

Centre for Epidemiology Versus Arthritis
University of Manchester

13/12/2022

Introduction

- Survival Analysis is concerned with the length of time before an event occurs.
- Initially, developed for events that can only occur once (e.g. death)
- Using time to event is more efficient that just whether or not the event has occured.
- It may be inconvenient to wait until the event occurs in all subjects.
- Need to include subjects whose time to event is not known (censored).

Plan of Talk

- Censoring
- Describing Survival
- Comparing Survival
- Modelling Survival

Censoring

- Exact time that event occured (or will occur) is unknown.
- Most commonly right-censored: we know the event has not occured yet.
- Maybe because the subject is lost to follow-up, or study is over.
- Makes no difference *provided* loss to follow-up is unrelated to outcome.

Censoring Examples: Chronological Time

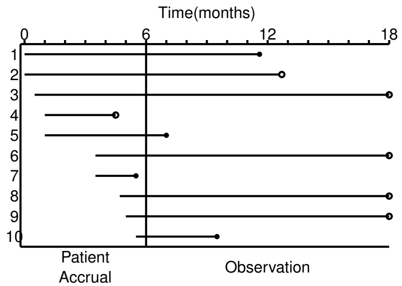

Censoring Examples: Followup Time

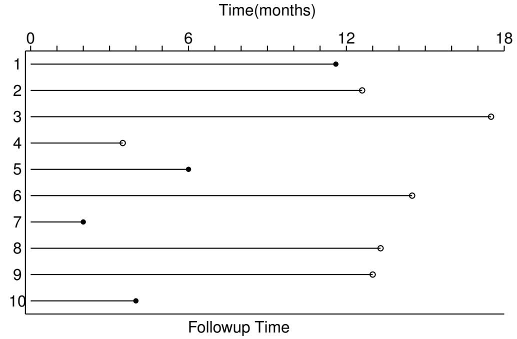

Other types of censoring

- Left Censoring:
  - Event had already occured before the study started.
  - Subject cannot be included in study.
  - May lead to bias.
- Interval Censoring:
  - We know event occured between two fixed times, but not exactly when.
  - E.g. Radiological damage: only picked up when film is taken.

Describing Survival: Survival Curves

- Survivor function: $S(t)$ probability of surviving to time t.
- If there are r_k subjects at risk during the k^{th} time-period, of whom f_k fail, probability of surviving this time-period for those who reach it is
$$\frac{r_k - f_k}{r_k}$$
- Probability of surviving the end of the k^{th} time-period is the probability of surviving to the end of the $(k-1)^{th}$ time-period, times the probability of surviving the k^{th} time-period. i.e

$$S(k) = S(k-1) \times \frac{r_k - f_k}{r_k}$$

Motion Sickness Study

- 21 subjects put in a cabin on a hydraulic piston,
- Bounced up and down for 2 hours, or until they vomited, whichever occured first.
- Time to vomiting is our survival time.
- Two subjects insisted on ending the experiment early, although they had not vomited (censored).
  - Is censoring independent of expected event time ?
- 14 subjects completed the 2 hours without vomiting.
- 5 subjects failed

Introduction
Censoring
Describing Survival
Comparing Survival
Modelling Survival
Survivor function
Stata Commands

Motion Sickness Study Life-Table

ID	Time	Censored	r_k	f_k	$S(t)$	
1	30	No	**21**	**1**	$\mathbf{20/21}$	$= \mathbf{0.952}$
2	50	No	**20**	**1**	$\mathbf{19/20 \times S(30)}$	$= \mathbf{0.905}$
3	50	Yes	**19**	**0**	$\mathbf{19/19 \times S(50)}$	$= \mathbf{0.905}$
4	51	No	**18**	**1**	$\mathbf{17/18 \times S(50)}$	$= \mathbf{0.855}$
5	66	Yes	**17**	**0**	$\mathbf{17/17 \times S(51)}$	$= \mathbf{0.855}$
6	82	No	**16**	**1**	$\mathbf{15/16 \times S(66)}$	$= \mathbf{0.801}$
7	92	No	**15**	**1**	$\mathbf{14/15 \times S(82)}$	$= \mathbf{0.748}$
8	120	Yes	**14**	**0**	$\mathbf{14/14 \times S(92)}$	$= \mathbf{0.748}$
⋮						
21	120	Yes	**14**	**0**	$\mathbf{14/14 \times S(92)}$	$= \mathbf{0.748}$

Introduction
Censoring
Describing Survival
Comparing Survival
Modelling Survival
Survivor function
Stata Commands

Kaplan Meier Survival Curves

- Plot of $S(t)$ against (t).
- Always start at (0, 1).
- Can only decrease.
- Drawn as a step function, with a downwards step at each failure time.

Introduction
Censoring
Describing Survival
Comparing Survival
Modelling Survival
Survivor function
Stata Commands

Stata commands for Survival Analysis

- `stset`: sets data as survival
  - Takes one variable: followup time
  - Option `failure` = 1 if event occurred, 0 if censored
- `sts list`: produces life table
- `sts graph`: produces Kaplan Meier plot

Stata Output

```
sts list if group == 1

         failure _d:  fail
   analysis time _t:  time

           Beg.          Net            Survivor      Std.
  Time    Total   Fail   Lost           Function     Error     [95% Conf. Int.]
-------------------------------------------------------------------------------
    30       21      1      0             0.9524    0.0465     0.7072    0.9932
    50       20      1      1             0.9048    0.0641     0.6700    0.9753
    51       18      1      0             0.8545    0.0778     0.6133    0.9507
    66       17      0      1             0.8545    0.0778     0.6133    0.9507
    82       16      1      0             0.8011    0.0894     0.5519    0.9206
    92       15      1      0             0.7477    0.0981     0.4946    0.8868
   120       14      0     14             0.7477    0.0981     0.4946    0.8868
-------------------------------------------------------------------------------
```


Kaplan Meier Curve: example

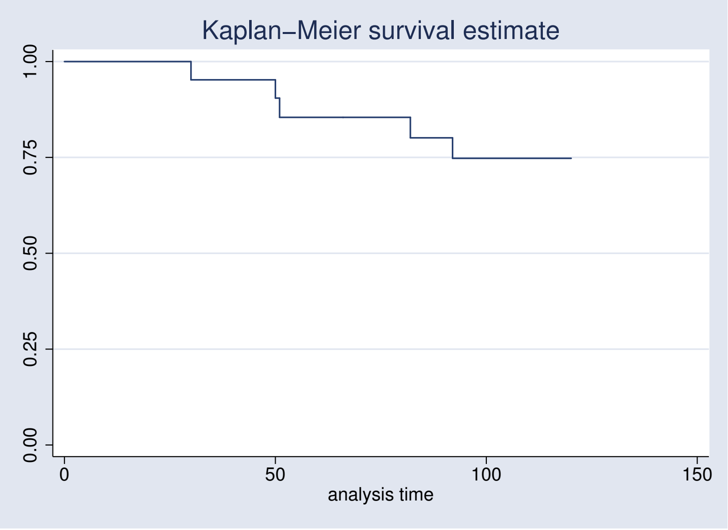

Comparing Survivor Functions

- **Null Hypothesis** Survival in both groups is the same
- **Alternative Hypothesis**
  1. Groups are different
  2. One group is consistently better
  3. One group is better at fixed time *t*
  4. Groups are the same until time *t*, one group is better after
  5. One group is worse up to time *t*, better afterwards.
- No test is equally powerful against all alternatives.

Comparing Survivor Functions

- Can use
  - Logrank test
    - Most powerful against consistent difference
  - Modified Wilcoxon Test
    - Most powerful against early differences
  - Regression
- Should decide which one to use beforehand.

Motion Sickness Revisited

- Less than 1/3 of subjects experienced an endpoint in first study.
- Further 28 subjects recruited
- Freqency and amplitude of vibration both doubled
- Intention was to induce vomiting sooner
- Were they successful ?

Comparing Survival Curves

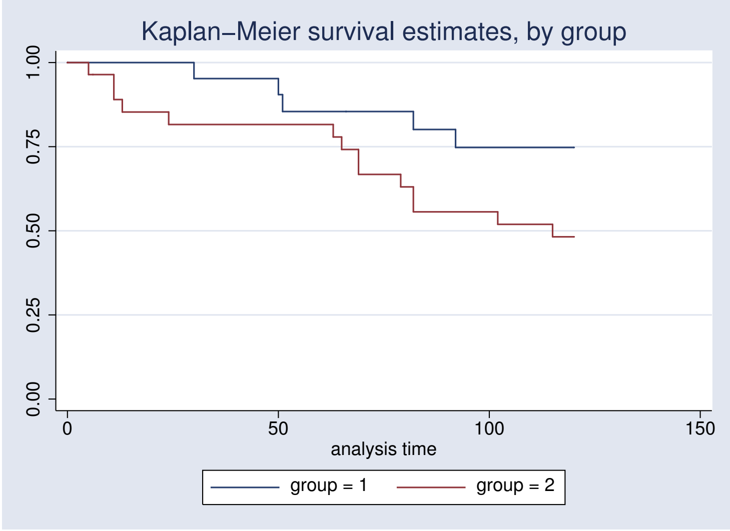

Comparison of Survivor Functions

- `sts test group` gives logrank test for differences between groups
- `sts test group, wilcoxon` gives Wilcoxon test

Test	χ^2	p
Logrank	3.21	0.073
Wilcoxon	3.18	0.075

What to avoid

- Compare mean survival in each group.
  - Censoring makes this meaningless
- Overinterpret the tail of a survival curve.
  - There are generally few subjects in tails
- Compare proportion surviving in each group at a fixed time.
  - Depends on arbitrary choice of time
  - Lacks power compared to survival analysis
  - Fine for description, not for hypothesis testing

Modelling Survival

- Cannot often simply compare groups, must adjust for other prognostic factors.
- Predicting survival function S is tricky.
- Easier to predict the hazard function.
  - Hazard function $h(t)$ is the risk of dying at time t, given that you've survived until then.
  - Can be calculated from the survival function.
  - Survival function can be calculated from the hazard function.
  - Hazard function easier to model

The Hazard Function

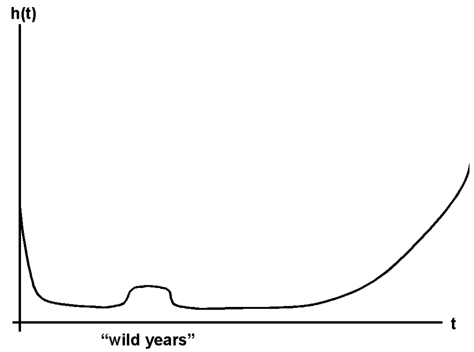

Hazard for all cause mortality for time since birth

Options for Modelling Hazard Function

- Parametric Model
- Semi-parametric models
  - Cox Regression (unrestricted baseline hazard)
  - Smoothed baseline hazard

Comparing Hazard Functions

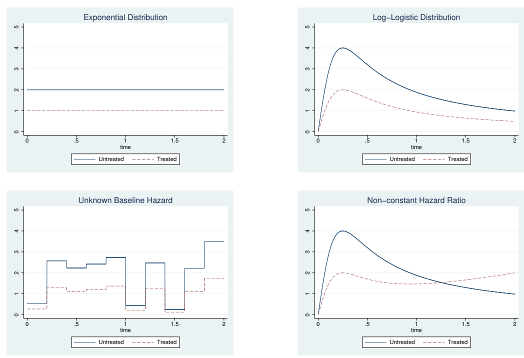

Parametric Regression

- Assumes that the shape of the hazard function is known.
- Estimates parameters that define the hazard function.
- Need to test that the hazard function is the correct shape.
- Was only option at one time.
- Now that semi-parametric regression is available, not used unless there are strong *a priori* grounds to assume a particular distribution.
- More powerful than semi-parametric if distribution is known

Cox (Proportional Hazards) Regression

- Assumes shape of hazard function is unknown
- Given covariates **x**, assumes that the hazard at time t,

$$h(t, x) = h_0(t) \times \Psi(\mathbf{x})$$

  where $\Psi = \exp(\beta_1 x_1 + \beta_2 x_2 + \ldots)$.
- Semi-parametric: h_0 is non-parametric, Ψ is parametric.
- t affects h_0, not Ψ
- **x** affects Ψ, not h_0

Cox Regression: Interpretation

Suppose x_1 increases from x_0 to $x_0 + 1$,

$$\begin{aligned} h(t, x_0) &= h_0(t) \times e^{(\beta_1 x_0)} \\ h(t, x_0 + 1) &= h_0(t) \times e^{(\beta_1 (x_0+1))} \\ &= h_0(t) \times e^{(\beta_1 x_0)} \times e^{\beta_1} \\ &= h(t, x_0) \times e^{\beta_1} \\ \Rightarrow \quad \frac{h(t, x_0+1)}{h(t, x_0)} &= e^{\beta_1} \end{aligned}$$

i.e. the **Hazard Ratio** is e^{β_1}

- Results may be presented as β or e^{β}
- $\beta > 0 \Rightarrow e^{\beta} > 1 \Rightarrow$ risk increased
- $\beta < 0 \Rightarrow e^{\beta} < 1 \Rightarrow$ risk decreased
- Should include a confidence interval.

Cox Regression: Testing Assumptions

- We assume hazard ratio is constant over time: should test.
- Possible tests:
  - Plot observed and predicted survival curves: should be similar.
  - Plot $-\log(-\log(S(t)))$ against $\log(t)$ for each group: should give parallel lines.
  - Formal statistical test:
    - Overall
    - Each variable
- May need to fit interaction between time period and predictor: assume constant hazard ratio on short intervals, not over entire period.

Cox Regression in Stata

- `stcox` *`varlist`* performs regression using *`varlist`* as predictors
- Option `nohr` gives coefficients in place of hazard ratios

Testing Proportional Hazards

- `stcoxkm` produced plots of observed and predicted survival curves
- `stphplot` produces $-\log(-\log(S(t)))$ against $\log(t)$ (log-log plot)
- `estat phtest` gives overall test of proportional hazards
- `estat phtest, detail` gives test of proportional hazards for each variable.

Cox Regression: Example

```
. stcox i.group

Cox regression -- Breslow method for ties

No. of subjects =           49                     Number of obs    =        49
No. of failures =           19
Time at risk    =         4457
                                                   LR chi2(1)       =      3.32
Log likelihood  =   -67.296458                     Prob > chi2      =    0.0685

------------------------------------------------------------------------------
          _t | Haz. Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     2.group |    2.45073   1.277744     1.72   0.086     .8820678    6.809087
------------------------------------------------------------------------------
```

Testing Assumptions: Kaplan-Meier Plot

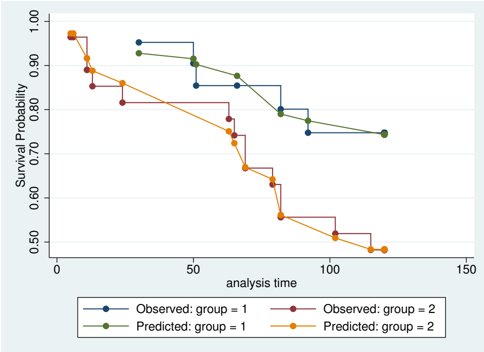

Testing Assumptions: log-log plot

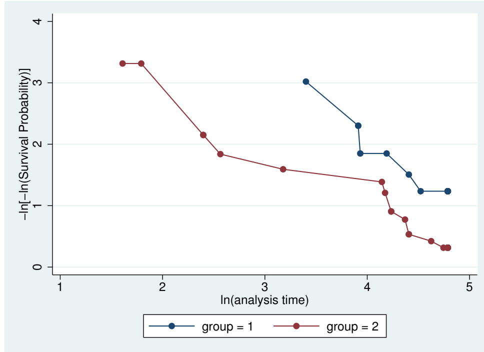

Testing Assumptions: Formal Test

```
. estat phtest

      Test of proportional hazards assumption
      ----------------------------------------------
                  |   chi2       df       Prob>chi2
      ------------+---------------------------------
      global test |   0.03        1         0.8585
      ----------------------------------------------
```


Allowing for Non-Proportional Hazards

- Effect of covariate varies with time
- Need to produce different estimates of effects at different times
- Use `stsplit` to split one record per person into several
- Fit covariate of interest in each time period separately

Non-Proportional Hazards Example

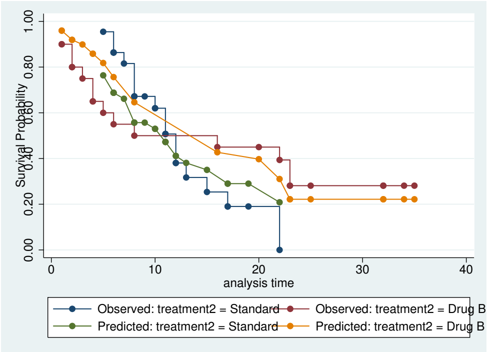

Non-Proportional Hazards Example

```
. stcox i.treatment2

------------------------------------------------------------------------------
          _t | Haz. Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  treatment2 |   .7462828   .3001652    -0.73   0.467     .3392646    1.641604
------------------------------------------------------------------------------

. estat phtest

      Test of proportional hazards assumption

      Time:  Time
      ----------------------------------------------------------------
                  |                   chi2       df       Prob>chi2
      ------------+---------------------------------------------------
      global test |                  10.28        1         0.0013
      ----------------------------------------------------------------
```


Non-Proportional Hazards Example: Fitting time-varying effect

```
stsplit period, at(10)
gen t1 = treatment2*(period == 0)
gen t2 = treatment2*(period == 10)

. stcox t1 t2

------------------------------------------------------------------------------
          _t | Haz. Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
          t1 |   1.836938   .8737408     1.28   0.201     .7231357    4.666262
          t2 |   .1020612   .0853529    -2.73   0.006     .0198156    .5256703
------------------------------------------------------------------------------

. estat phtest

      Test of proportional hazards assumption

      Time:  Time
      ----------------------------------------------------------------
                  |                   chi2       df       Prob>chi2
      ------------+---------------------------------------------------
      global test |                   1.34        2         0.5114
      ----------------------------------------------------------------
```


Non-Proportional Hazards Example

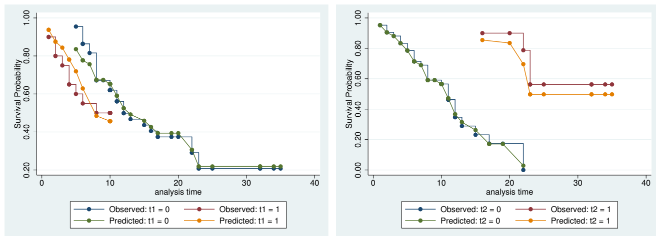

Time varying covariates

- Normally, survival predicted by baseline covariates
- Covariates may change over time
- Can have several records for each subject, with different covariates
- Each record ends with a censoring event, unless the event of interest occurred at that time
- Need to have unique identifier for each individual so that stata knows which observations belong together
- Option `tvc()` is for variables that increase linearly with time