Survival analysis – two groups

Theis Lange
Department of Biostatistics, University of Copenhagen
&
Center for Statistical Science, Peking University.

Mail: thlan@sund.ku.dk

Survival data

The outcome is the *time to* a specific event happens.
Examples:

- Time from treatment start to death.
- Time from birth to death.
- Time from employment to first sick leave.
- Time from pregnancy to birth.

Survival data is characterized by:

- Right skew (typically).
- Only partially observed.

In conclusion: *Regular methods cannot be employed.*

A reading suggestion:

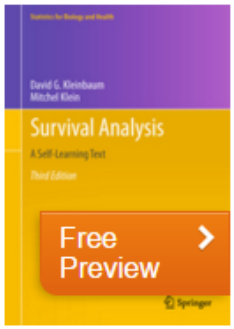

Survival Analysis

A Self-Learning Text, Third Edition

Authors: **Kleinbaum**, David G., **Klein**, Mitchel

Survival function

- A survival time is just a numerical variable, but means and standard deviations are not good descriptions.

- Instead focus on the survival function.
  With T denoting a survival time the survival function is
  $S(t) = P(T > t)$
  = "Probability of being alive at time t"

- The survival function satisfies:
  - S(t) ≥ 0 for all t
  - Non-increasing.
  - $S(-\infty) = 1$ (typically $S(0) = 1$)
  - $S(\infty) = 0$

- If we knew survival function we knew all relevant information.

Hazard function

- The survival function describes the whole life-time distribution.

- If here-and-now measure is wanted use instead the hazard-function defined as

$$h(t) \approx \frac{P(t \leq T < t + d \mid T \geq t)}{d}$$

for *d* tending to zero.

- Hazard value can be thought of as the probability of dying within the next 1 year (or whatever units we are using for the time scale).

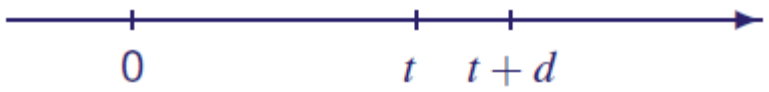

Survival function vs. hazard function

- Think of life as driving along in a car

- Survival function is the odometer

- Hazard function is the speedometer.

Censoring

Survival data is often *right-censored*. This implies that you only know a lower limit for the event time of interest.

Examples:

- The study ends before all participants have died.
- We loose track of the patient during follow-up (immigration etc.)

Data can also be left censored when only an upper limit is known.

- Time to HIV infection vs. time to first positive HIV test.
- Age a child learn to read vs. time to positive reading test.

Remission time for acute Leukemia

- Example from

  Freirich et al. (1963). The effect of 6-mercaptopurine on the duration of remission time of steroid induced remission in acute leukaemia. *Blood*, **21** 699:716.

- 42 patients randomized to either placebo or 6-MP treatment.

- Patients included in the period 1959 to 1960.

Remission data

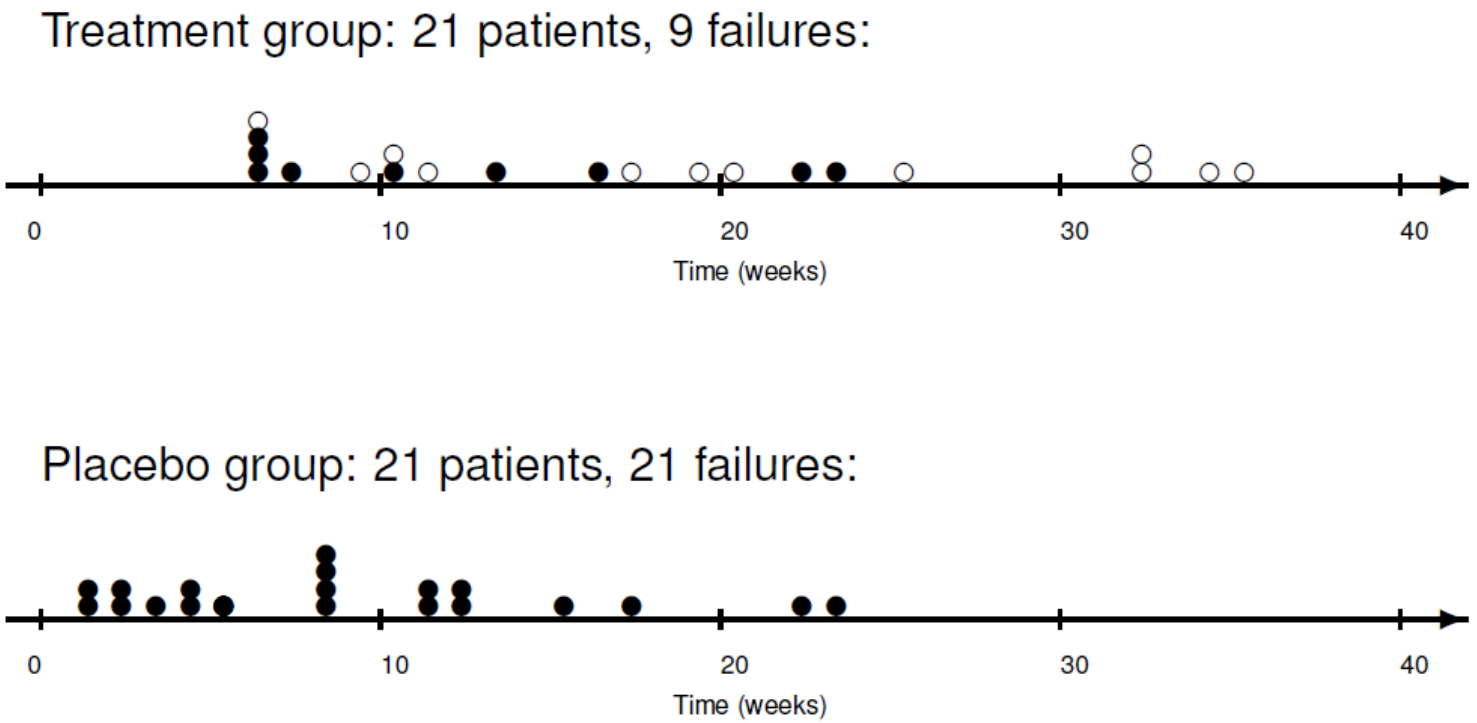

Full circle denotes death.

Empty circle denotes censoring.

Remission data

The raw data will look like this

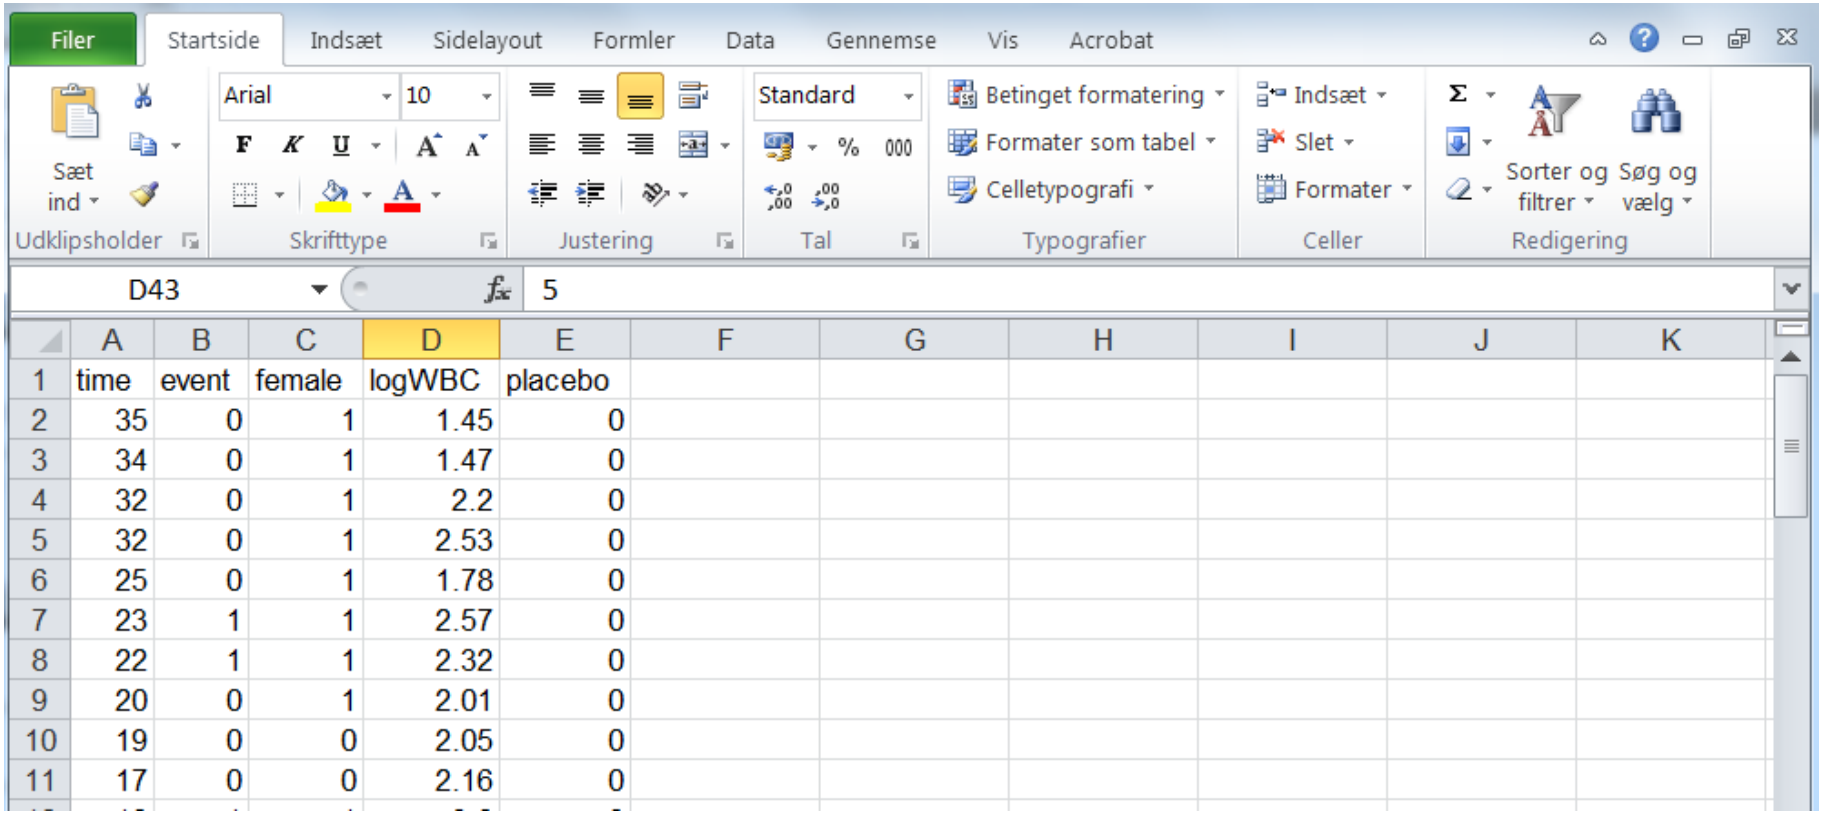

	A	B	C	D	E	F	G	H	I	J	K
1	time	event	female	logWBC	placebo						
2	35	0	1	1.45	0						
3	34	0	1	1.47	0						
4	32	0	1	2.2	0						
5	32	0	1	2.53	0						
6	25	0	1	1.78	0						
7	23	1	1	2.57	0						
8	22	1	1	2.32	0						
9	20	0	1	2.01	0						
10	19	0	0	2.05	0						
11	17	0	0	2.16	0						

The variable “event” is 1 for deaths and 0 for censoring.

Estimating the survival function

- We cannot estimate survival as #alive/#in total – **why not?**
- Use instead Kaplan-Meier procedure.
- In treatment group we have
  No. still under risk

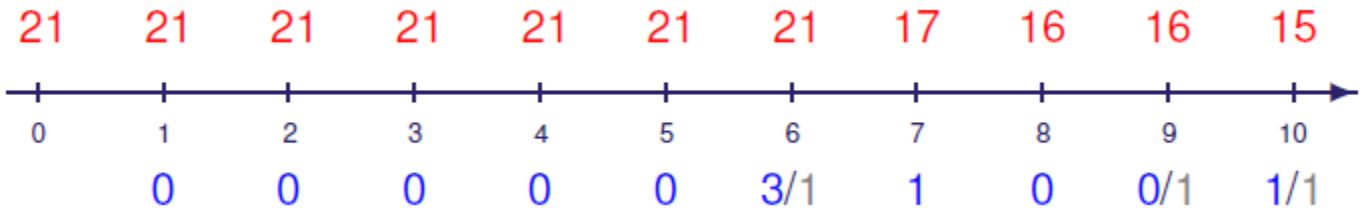

Number of events/number of censorings

- Survival function is then estimated by

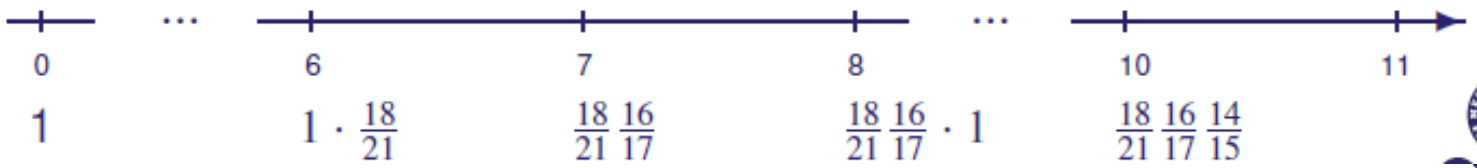

Kaplan-Meier plot for treatment group

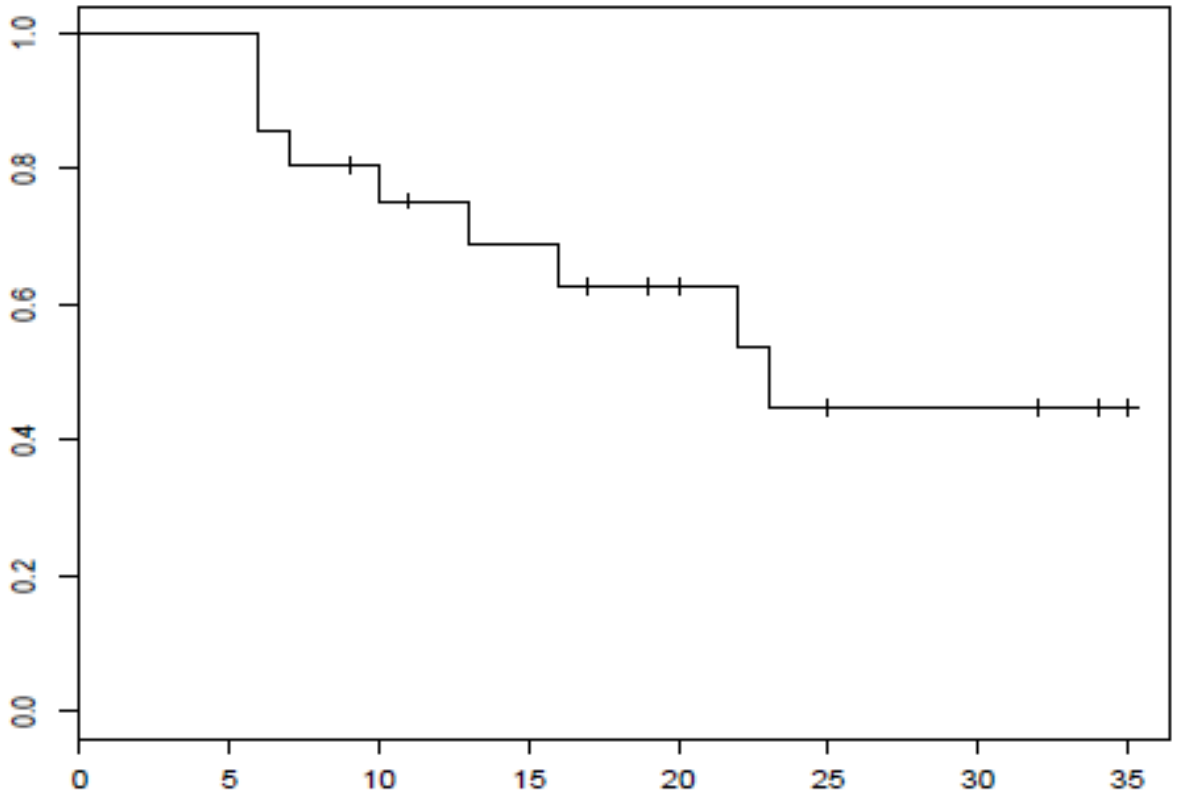

What is median survival time?

KM plot for treatment group with confidence interval

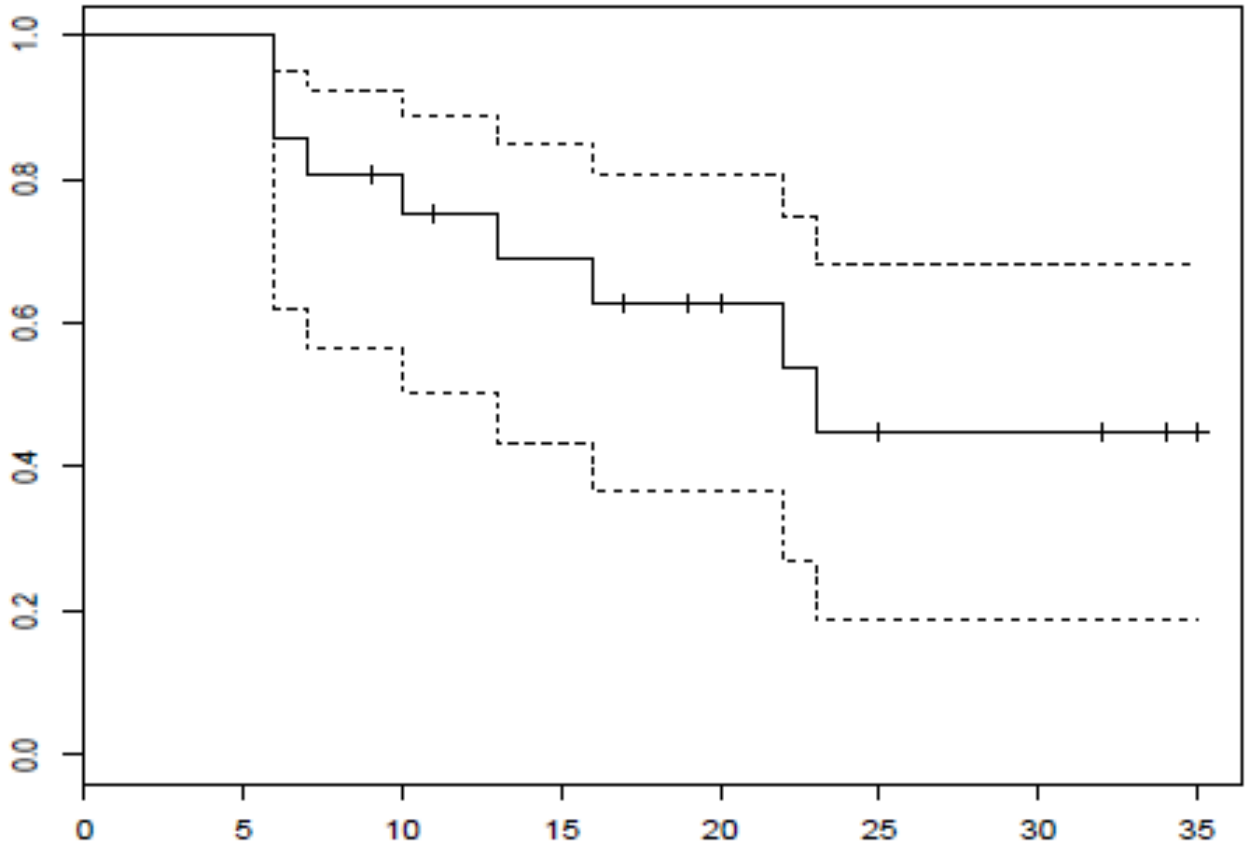

What is 6 month survival probability?

And with which precision is this known?

KM plot for both treatment groups

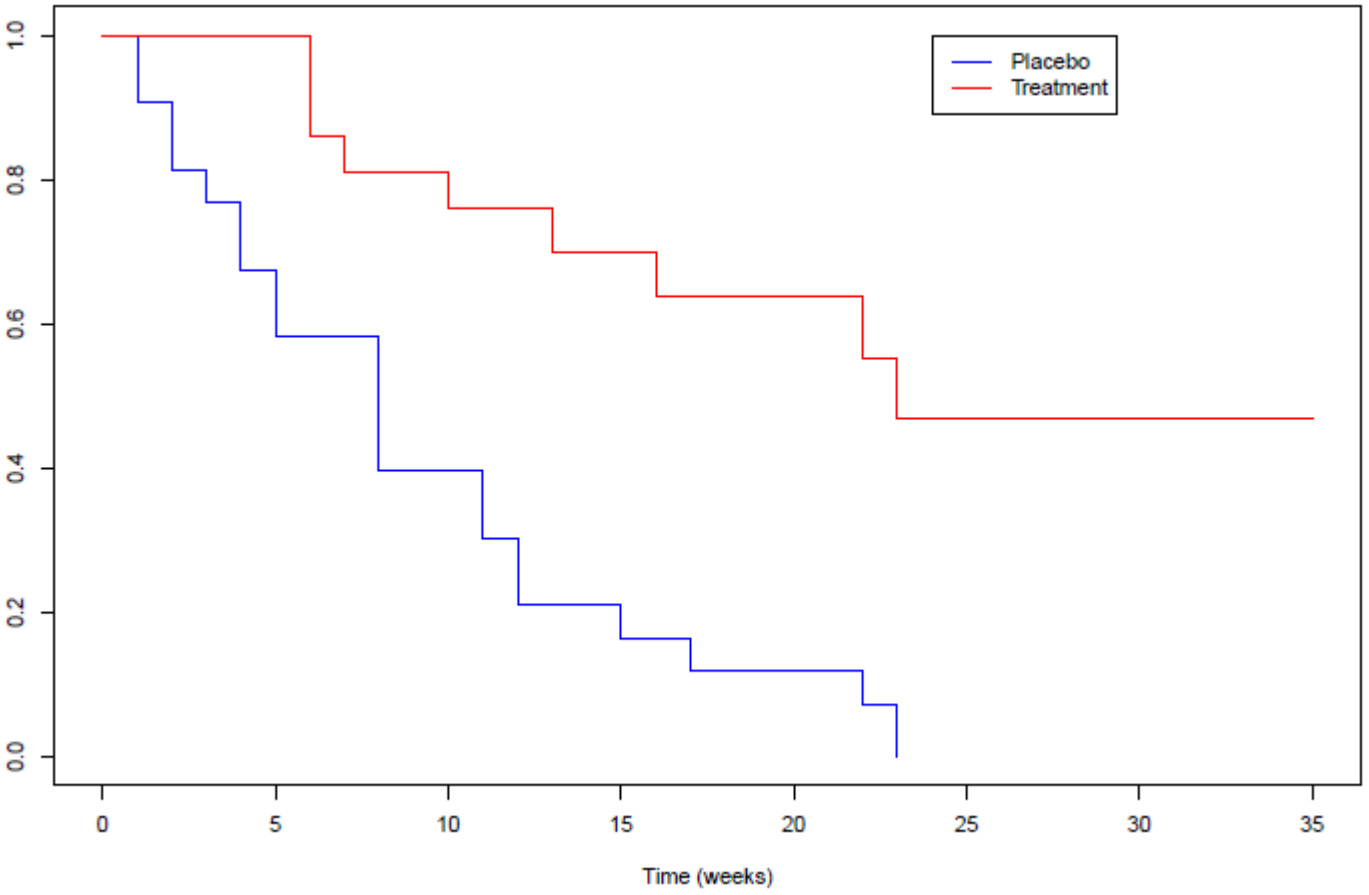

Software: Getting KM plots in R

```
remisData <- read.csv2("remissionData.csv")
head(remisData)
library(survival)

# to get plot with both curves
survFitObj1 <- survfit(Surv(time, event)~placebo,
   data=remisData, conf.int = 0.95)
plot(survFitObj1)

# to get plot with one curve AND confidence bands
survFitObj2 <- survfit(Surv(time, event)~1,
   data=subset(remisData, placebo==0), conf.int = 0.95)
plot(survFitObj2)
```

Getting life tables in R

```
> survFitObj2 <- survfit(Surv(time, event)~1, data=subset(remisData,
   placebo==0), conf.int = 0.95)
> summary(survFitObj2)
Call: survfit(formula = Surv(time, event) ~ 1, data =
   subset(remisData,
    placebo == 0), conf.int = 0.95)

 time n.risk n.event survival std.err lower 95% CI upper 95% CI
    6     21       3    0.857  0.0764        0.720        1.000
    7     17       1    0.807  0.0869        0.653        0.996
   10     15       1    0.753  0.0963        0.586        0.968
   13     12       1    0.690  0.1068        0.510        0.935
   16     11       1    0.627  0.1141        0.439        0.896
   22      7       1    0.538  0.1282        0.337        0.858
   23      6       1    0.448  0.1346        0.249        0.807
>
```

Assumptions about censoring

- Survival analysis builds on the assumptions of *non-informative censoring*.

- In Kleinbaum and Klein this is described as

  > **Non-informative censoring** occurs if the distribution of survival times (T) provides no information about the distribution of censorship times (C), and vice versa.

- You can also loosely think of it as: "Would knowing censoring as happened help you to predict event time?". If the answer is yes, you have a problem.

- You cannot not formally test for non-informative censoring.

- Exercise: Can you censor people because treatment was stopped due to severe side effects?

Comparing two survival curves

- The standard non-parametric tool for comparing two KM-curves is the log-rank test.

- The underlying idea is similar to chi-sq test.
  A. Compute expected event counts assuming no difference in survival functions (ie. by treating the whole sample as one group).
  B. Compare these to the observed counts in one of the groups.
  C. Add together over all event times in that group.

$$E_{1j} = \frac{O_j}{N_j} N_{1j}$$

- Note that as in the KM-plots censoring is handled by every time conditioning on the number still at risk.

The log-rank test in R

```
> survdiff(Surv(time, event)~placebo, data=remisData)

Call:
survdiff(formula = Surv(time, event) ~ placebo, data =
   remisData)

            N Observed Expected (O-E)^2/E (O-E)^2/V
placebo=0  21        9     19.3      5.46      16.8
placebo=1  21       21     10.7      9.77      16.8

 Chisq= 16.8  on 1 degrees of freedom, p= 4.17e-05
```


The log-rank test in R (with exact p-value)

```
> library(coin)
> logrank_test(Surv(time, event)~factor(placebo),
   data=remisData, distribution = "exact")

   Exact Two-Sample Logrank Test

data:  Surv(time, event) by factor(placebo) (0, 1)
Z = 3.9034, p-value = 2.612e-05
```


Conclusion

- Survival data often have non-symmetric distributions.
- However, it is the presence of censoring that makes standard methods invalid.
- Censoring must be non-informative for (standard) survival analyses tools to work.
- Survival function is estimated by Kaplan-Meier plots.
- Comparisons (ie. p-values) are made using log-rank test.

A case: The 6S trial

6S - **S**candinavian **S**tarch for **S**evere **S**epsis/**S**eptic **S**hock trial

- **Background:**
  Intravenous fluids are the mainstay of treatment for patients with hypovolemia due to severe sepsis to obtain fast circulatory stabilisation.
  Commonly applied interventions:
  - Crystalloids including saline and dextrose (Ringer's solution)
  - Colloids containing larger molecules such as starch or gelatine.
- **Preferred choice in Scandinavian intensive care units (ICU)**
  Hydtroxyethyl starch (HES) 130/0.42
  (previously high molecular weight HES - reported to cause acute kidney failure and bleeding)
- **However**
  HES 130/0.42 is largely unstudied in patients with severe sepsis;
  Lack of efficacy data and concerns about safety

6S – the details

- **Study population:**
  Patients with severe sepsis admitted to an intensive care unit (ICU) who needed fluid for circulatory stabilisation
- **Randomisation to:**
  - Hydroxyethyl starch (HES 130/0.42)
  - Ringer's solution
- **Primary outcome:**
  Death or end-stage kidney failure at 90 days after randomisation

➢ In total was 798 patients included.

The NEW ENGLAND JOURNAL of MEDICINE

ORIGINAL ARTICLE

Hydroxyethyl Starch 130/0.42 versus Ringer's Acetate in Severe Sepsis

6S – baseline data

Table 1. Baseline Characteristics of the Patients.*

Characteristic	HES 130/0.42 (N = 398)	Ringer's Acetate (N = 400)
Age — yr		
Median	66	67
Interquartile range	56–75	56–76
Male sex — no. (%)	239 (60)	244 (61)
Ideal body weight — kg†		
Median	72	72
Interquartile range	60–80	60–80
Admitted to university hospital — no. (%)	194 (49)	188 (47)
Surgery — no. (%)‡		
Emergency	114 (29)	116 (29)
Elective	34 (9)	48 (12)
Source of ICU admission — no. (%)		
Emergency department	109 (27)	94 (24)
General ward	177 (44)	196 (49)
Operating or recovery room	59 (15)	54 (14)
Other ICU in the same hospital	21 (5)	14 (4)
Other hospital	32 (8)	42 (10)
Source of sepsis — no. (%)§		

6S – treatment received

Table 2. Fluid Therapy before and after Randomization.*

Variable	HES 130/0.42 (N = 398)			Ringer's Acetate (N = 400)			P Value†
	Patients	Volume Received‡		Patients	Volume Received‡		
		median	interquartile range		median	interquartile range	
	no./total no.§	*ml*		*no./total no.§*	*ml*		
Trial fluid							
Day 1¶	374/397	1500	1000–1500	375/400	1500	1000–2000	0.09
Day 2	288/379	1500	1000–2000	307/380	1500	950–2000	0.50
Day 3	176/330	1000	500–1500	170/326	1000	500–1500	0.78

6S – outcomes analyzed by 2 by 2 technique

Outcome	HES 130/0.42 (N = 398)	Ringer's Acetate (N = 400)
Primary outcome		
Dead or dependent on dialysis at day 90 — no. (%)	202 (51)	173 (43)
Dead at day 90 — no. (%)	201 (51)	172 (43)
Dependent on dialysis at day 90 — no. (%)	1 (0.25)	1 (0.25)

Or in R:

```
chisq.test(kidneyData$mortality_90days, kidneyData$intervention,
   correct = F)

library(epitools)
riskratio(kidneyData$intervention, kidneyData$mortality_90days)
```

6S – Survival analysis

- Why so few censorings?

- In R:

```
fit1 <- survfit(Surv(time_to_death, mortality_90days) ~ intervention, data=kidneyData)
plot(fit1)
```

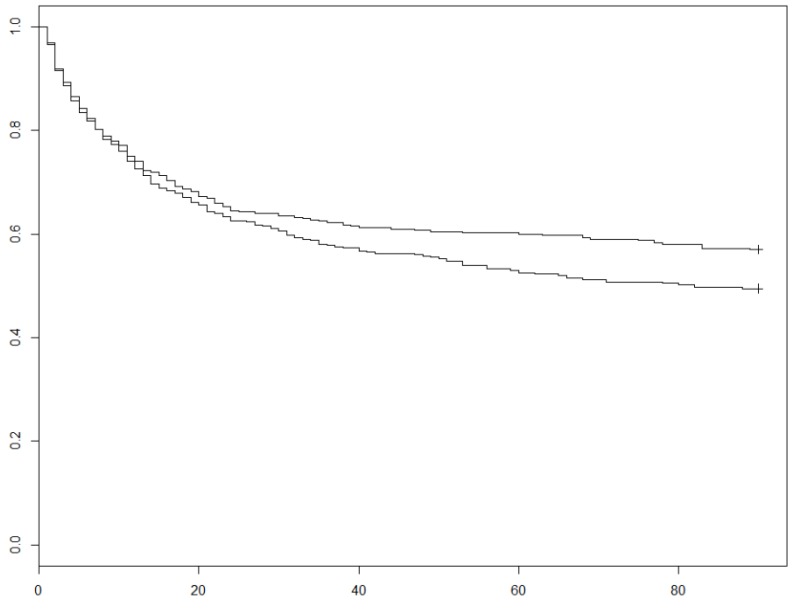

6S – Survival analysis

- Or a bit nicer

```
fit1 <- survfit(Surv(time_to_death, mortality_90days) ~ intervention,
        data=kidneyData)
plot(fit1, col=1:2, ylim=c(0.5, 1), xlab="Days", ylab="Survival pct.")
legend("topright", legend=c("HES 130/0.42", "Ringer's acetate"), lty=1,
        col=1:2, inset = 0.05)
```

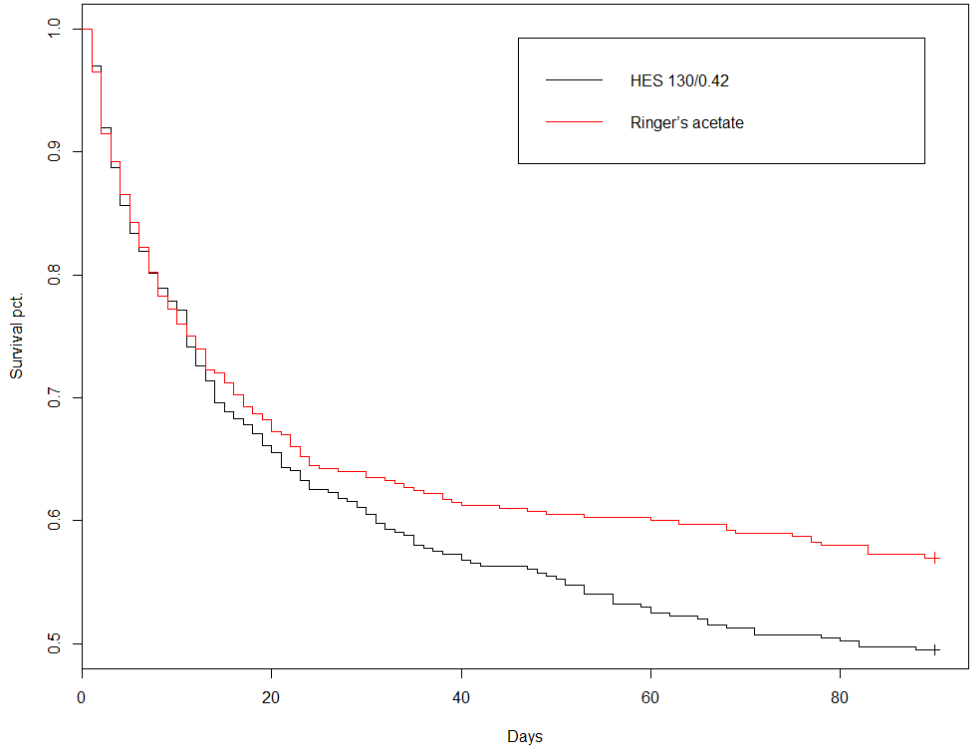

6S – The published plot

A Time to Death

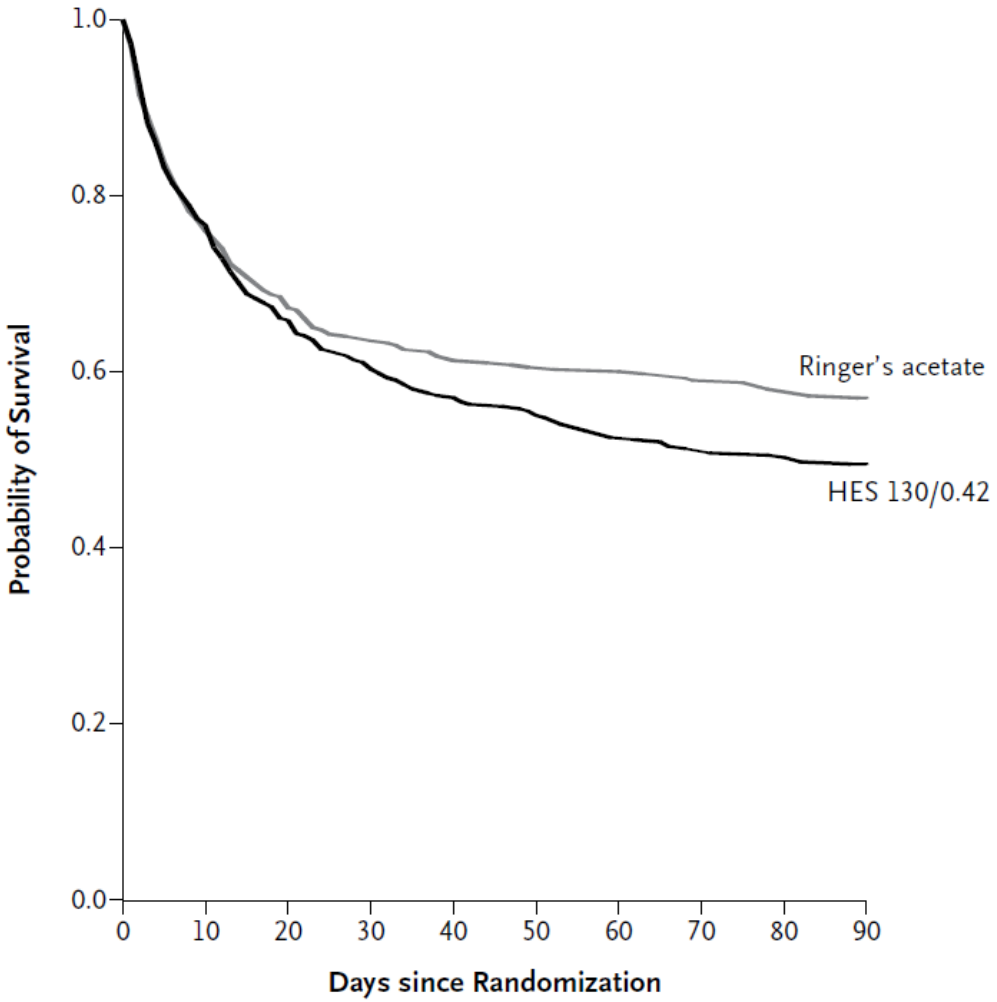

No. at Risk

HES 130/0.42	398	240	209	197
Ringer's acetate	400	254	240	228

6S – the log rank test

```
> survdiff(Surv(time_to_death,
   mortality_90days)~intervention, data=kidneyData)
```

```
Call:
survdiff(formula = Surv(time_to_death, mortality_90days) ~
   intervention,
    data = kidneyData)

                N Observed Expected (O-E)^2/E (O-E)^2/V
intervention=0 398      201      184      1.65      3.31
intervention=1 400      172      189      1.60      3.31

 Chisq= 3.3  on 1 degrees of freedom, p= 0.069
```