Survival analysis:
Competing risks

Mark Stevenson

Faculty of Veterinary and Agricultural Sciences

The University of Melbourne, Parkville Victoria 3010 Australia

mark.stevenson1@unimelb.edu.au

Roadmap

- Introduction
- A worked example

Introduction

- So far, we've described analytical approaches that can be used when the outcome interest is the same for all study subjects
  - e.g. when calving to conception interval is the outcome of interest, we monitor time to a single event for all subjects, conception

- Sometimes, it we might want to distinguish between several different kinds of events

Introduction

- Factors influencing length of productive life in dairy cows
  - monitor a population for a defined period
  - record both the date and timing of removal from the herd
  - record the reason for removal

- The hazard of removal for reproductive failure might be quite different to the hazard of removal for udder disorders

Descriptive epidemiological study on culling and deaths in eight dairy herds

MA STEVENSON[a] and IJ LEAN
Department of Animal Science, University of Sydney, Camden, New South Wales 2570

Objectives To describe the production and reproduction characteristics of the population under investigation and to compare them with the target population of non-seasonally calving dairy herds, to describe the pattern of culling with respect to length of productive life and length of time after calving, and to identify stages of the production cycle that are

cow is incapable of further production) or economic (meaning that the cow is removed because a replacement is expected to produce greater profit). This classification is thought to promote a planned approach to culling, where the income-generating capacity of members of the herd are compared to that of potential herd replacements. Optimum herd prof-

Stevenson, M., Lean, I., 1998. Descriptive epidemiological study on culling and deaths in eight dairy herds. Australian Veterinary Journal 76, 482 - 488.

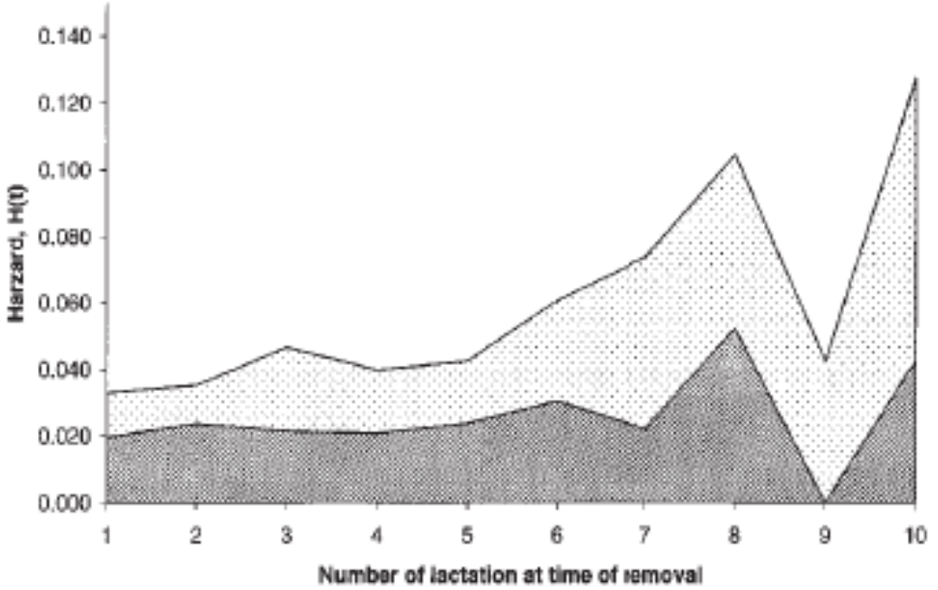

Figure 3. Cows culled for reproductive failure, hazard function for lactations commenced at time of removal. Dark shading depicts the hazard for removal for failure to conceive. Light shading depicts the hazard for other reproductive tract disorders. In this figure (and in Figure 4) the hazards are additive – the hazard of removal for failure to conceive plus the hazard for removal for other reproductive tract disorders at a given point in time represents the total hazard of removal for reproductive failure. The data presented here and in Figure 4 are pooled for all

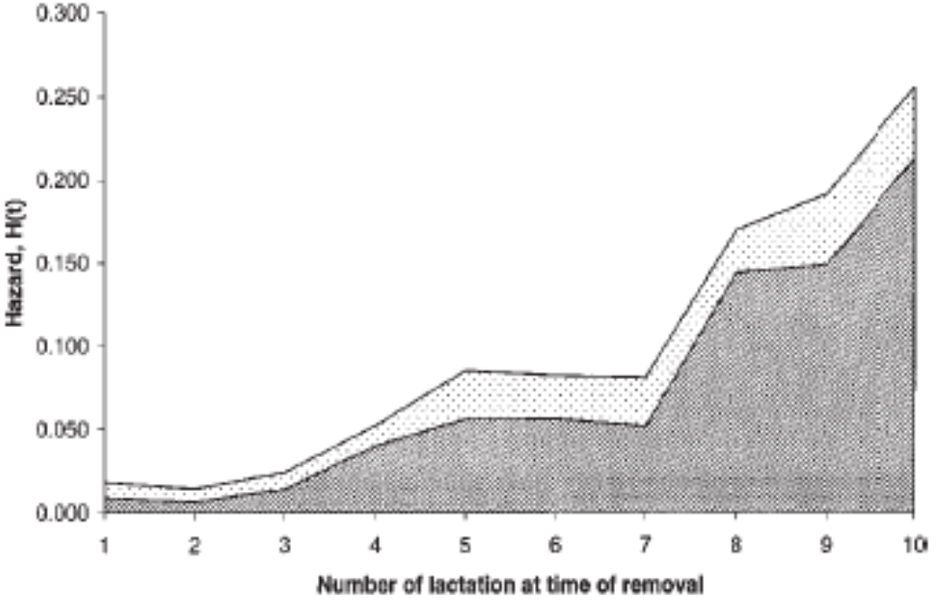

Figure 5. Cows culled for udder disorders, hazard function for lactations commenced at time of removal. Dark shading depicts the hazard of removal for acute and chronic mastitis (defined in Table 1). Light shading depicts the hazard for other udder disorders (defined in Table 1). In this figure (and in Figure 6) the hazards are additive – the hazard of removal for acute and chronic mastitis plus the hazard for removal for other udder disorders at a given point in time represents the total hazard of removal for other udder disorders. The data presented here and in Figure 6 are pooled for all herds.

Stevenson, M., Lean, I., 1998. Descriptive epidemiological study on culling and deaths in eight dairy herds. Australian Veterinary Journal 76, 482 - 488.

Introduction

- Key point: the occurrence of one type of event removes the individual from being at risk of all other event types
  - cows removed for reproductive failure are no longer at risk of removal for udder disorders
  - people who die from cancer are no longer at risk of death from heart disease
  - employees who resign are no longer at risk of being sacked

Introduction

- Let's assume that the events we are interested in removal reasons in dairy cattle and we've classified them into k = 5 types:
  - reproductive failure, udder disorders, lameness, low production and miscellaneous reasons

Introduction

- For each removal reason we define a separate hazard function which will be called a type-specific or cause-specific hazard
  - let t_i be a random variable denoting the time of removal for cow i
  - let J_k be a variable denoting removal reason (where k = 1, 2, 3, 4, and 5)

Introduction

- We can estimate a hazard function for removal for udder disorders (J_2), and in doing so we treat removals for all other reasons as censored observations
  - this gives us a reason-specific survival function, giving the probability that an event of type ($J_1, J_3, J_4,$ or J_5) occurs later than time t
  - now that removal type-specific hazards are defined we can proceed to formulate models for their dependence on covariates
  - both proportional hazards and accelerated failure time models can be fitted

Introduction

- For example, we can specify a general proportional hazards model for all five different removal reasons:

$$h(t, X) = hJ_{0(}t) \exp(\beta_1 x_{1i} + \beta_2 x_{2i} + \ldots \beta_m x_{mi})$$

- The coefficients β can be subscripted according to removal reason to quantify the effect of each covariate on each of the removal reasons

Introduction

- Some of the coefficients may be set to zero, excluding the influence of that covariate for a specific removal reason

- The baseline survival function $h_{J0}(t)$ is subscripted by removal reason to allow the dependence of the hazard on time to vary across removal reasons

Introduction

- Although it'd be unusual, there'd be nothing to stop you from specifying a Weibull model for reproductive failure, a gamma model for udder disorders and a proportional hazards model for all other removal types

- What makes this possible is that the models may be estimated separately for each event type, with no loss of statistical precision
- This is the most important principle of competing risks analysis

Introduction

- As well, you don't need to estimate models for all event types unless you really need to

- For example, if all you are really interested in is to identify risk factors for removal due to reproductive failure, then develop a single model for reproductive failure, treating all other removals as censored observations

Roadmap

- Introduction
- A worked example

A worked example

- Kyle (1993) studied 241 cases of monoclonal gammopathy identified before 1 January 1971 with between 20 and 35 years of total follow up on each patient
  - the response variable is the time to the first of various adverse events
  - death (n = 130), multiple myeloma (n = 39) and 'other' (n = 20)
- In a competing risks analysis, we assign one stratum for each outcome type and all subjects appear in each stratum

A worked example

- In the following example data set the first subject experiences death at day 760 and the second subject experiences lymphoproliferative disease at day 2160

```
library(survival); setwd("D:\\TEMP")
dat <- read.table("mgus2.csv", header = TRUE, sep = ",")
```

id	time	status	endpoint	sex	age	hgb	creat	mspike
1	760	1	death	2	79	1.5	1.2	2.0
1	760	0	myeloma	2	79	1.5	1.2	2.0
1	760	0	other	2	79	1.5	1.2	2.0
2	2160	0	death	2	76	13.3	1.0	1.8
2	2160	0	myeloma	2	76	13.3	1.0	1.8
2	2160	1	other	2	76	13.3	1.0	1.8

id: patient identifier.
time: day of event.
status: 0 = censored, 1 = died.
endpoint: reason for failure - death, myeloma or other lymphoproliferative disorder.
sex: 1 = male, 2 = female.
age: patient age at enrolment.
hgb: plasma haemoglobin concentration at enrolment.
creat: plasma creatinine concentration at enrolment.
mspike: size of monclonal spike.

```
mgus2.km <- survfit(Surv(time, status) ~ endpoint, type = "kaplan-
meier", data = dat)


plot(mgus2.km, xlab = "Days to endpoint", ylab = "Cumulative
proportion to experience event", lty = c(1,2,3), mark.time = FALSE)
legend(x = "topright", legend = c("Death","Myeloma","Other"), lty =
c(1,2,3), bty = "n")
```

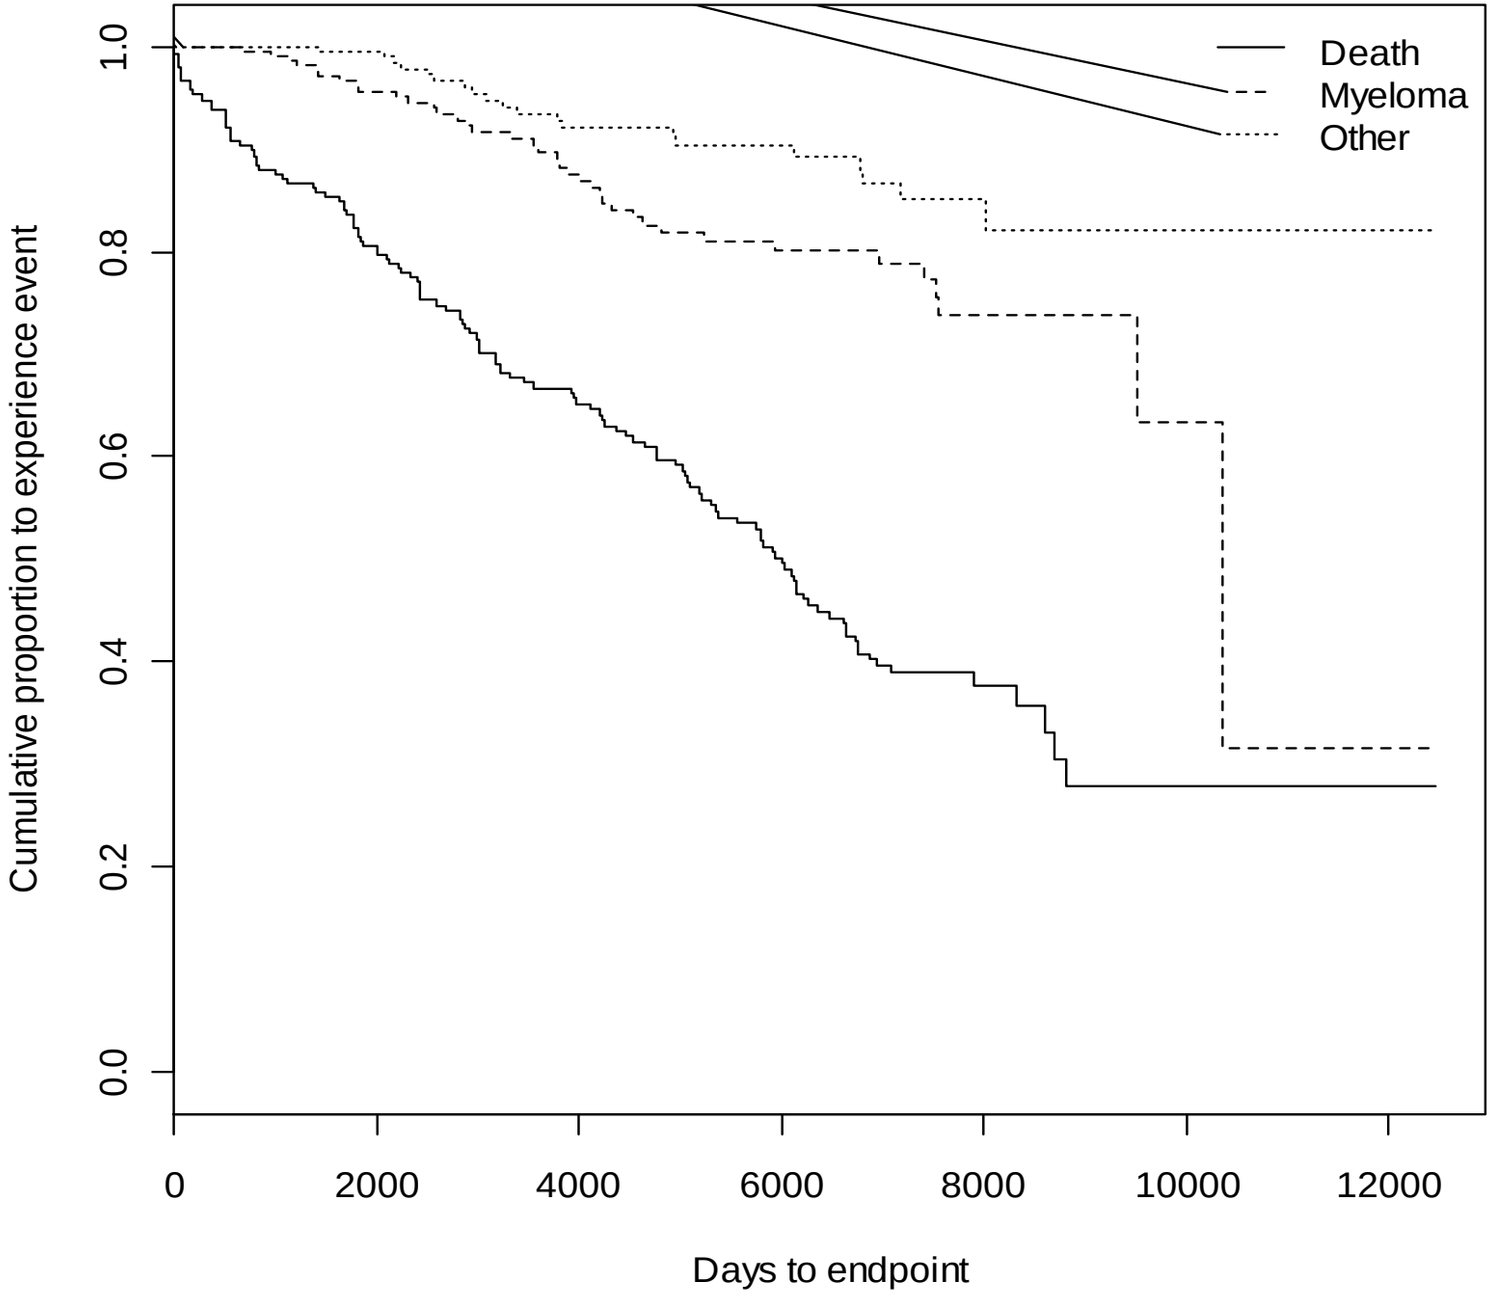
Death
Myeloma
Other
Cumulative proportion to experience event
0.0
0.2
0.4
0.6
0.8
1.0
0
2000
4000
6000
8000
10000
12000
Days to endpoint

A worked example

- Even though the data set contains three observations for each subject, because the endpoint times are independent it's OK to apply a log rank test to compare survivorship for the three outcomes

```
survdiff(Surv(time, status) ~ endpoint, data = dat, na.action =
na.omit, rho = 0)

Call:
survdiff(formula = Surv(time, status) ~ endpoint, data = dat,
    na.action = na.omit, rho = 0)

n=720, 3 observations deleted due to missingness.

                   N Observed Expected (O-E)^2/E (O-E)^2/V
endpoint=death   240      130     62.7      72.3     108.5
endpoint=myeloma 240       38     62.7       9.7      14.6
endpoint=other   240       20     62.7      29.0      43.6

 Chisq= 111  on 2 degrees of freedom, p= 0
```

Now fit a Cox model (after going through the usual process of screening covariates for their effect on time to event):

```
mgus2.cph01 <- coxph(Surv(time, status, type = "right") ~ sex + age
+ hgb + mspike + strata(endpoint), data = dat)
```

```
summary(mgus2.cph01)

Call:
coxph(formula = Surv(time, status, type = "right") ~ sex + age +
    hgb + mspike + cluster(id) + strata(endpoint), data = dat)

  n=708 (15 observations deleted due to missingness)

             coef exp(coef)  se(coef) robust se      z Pr(>|z|)
sex    -0.339901  0.711840  0.156319  0.152916 -2.223   0.0262 *
age     0.051373  1.052715  0.007350  0.007298  7.039 1.93e-12 **
hgb    -0.166434  0.846678  0.044324  0.035115 -4.740 2.14e-06 **
mspike -0.087751  0.915989  0.186866  0.182741 -0.480   0.6311
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

       exp(coef) exp(-coef) lower .95 upper .95
sex       0.7118      1.405    0.5275    0.9606
age       1.0527      0.950    1.0378    1.0679
hgb       0.8467      1.181    0.7904    0.9070
mspike    0.9160      1.092    0.6402    1.3105
```

The advantage of a large data set is that it allows for easy estimation of within-event-type coefficients. For instance, one might ask if the effect of age is identical for both outcomes, while controlling for the common effect of haemoglobin. This can be investigated by coding two dummy variables:

```
age1 <- dat$age * (dat$endpoint == "death")
age2 <- dat$age * (dat$endpoint != "death")

mgus2.cph02 <- coxph(Surv(time, status, type = "right") ~ sex + age1
+ age2 + hgb + mspike + strata(endpoint), data = dat)
summary(mgus2.cph02)
```

```
summary(mgus2.cph02)

Call:
coxph(formula = Surv(time, status, type = "right") ~ sex + age1 +
    age2 + hgb + mspike + strata(endpoint), data = dat)

  n=708 (15 observations deleted due to missingness)

            coef exp(coef)  se(coef)      z Pr(>|z|)
sex     -0.323615  0.723529  0.156364 -2.070 0.038487 *
age1     0.075578  1.078508  0.009315
age2     0.002642  1.002646  0.012252
hgb     -0.162255  0.850224  0.044615
mspike -0.088710  0.915111  0.187009
---
Signif. codes:  0 '***' 0.001 '**' 0.

       exp(coef) exp(-coef) lower .95 upper .95
sex       0.7235     1.3821    0.5325     0.983
age1      1.0785     0.9272    1.0590     1.098
age2      1.0026     0.9974    0.9789     1.027
hgb       0.8502     1.1762    0.7790     0.928
mspike    0.9151     1.0928    0.6343     1.320
```

Age is a significant predictor of death as an endpoint.

Age is, however, of less importance in predicting time to onset of plasma cell malignancy (i.e. myeloma or other lymphoproliferative disorders).

Roadmap

- Introduction
- A worked example