‘survextrap‘: flexible and transparent Bayesian survival modelling using combinations of data and judgements

Christopher Jackson
MRC Biostatistics Unit, University of Cambridge
with Iain Timmins (AstraZeneca), Michael Sweeting (GSK)

PSI Conference, 9 June 2025

Survival extrapolation: long-term decisions from short-term data

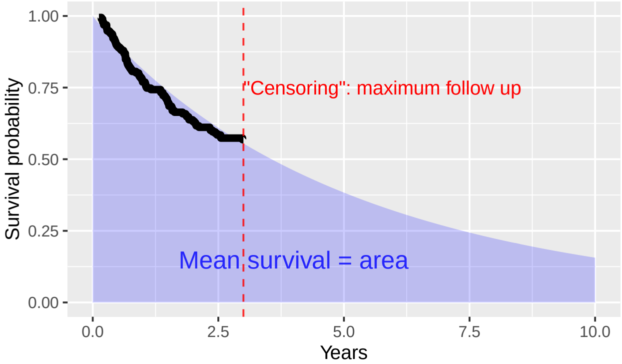

Limited trial follow up

Health economic decision making needs expected survival over the longer term

Area under the survival curve.
Estimated by parametric models...

Consequences of policy decisions will last longer than the end of the data

Parametric survival models: examples

Weibull distribution

Survivor function $S(t|\lambda, \alpha) = \exp(-\lambda t^{\alpha})$

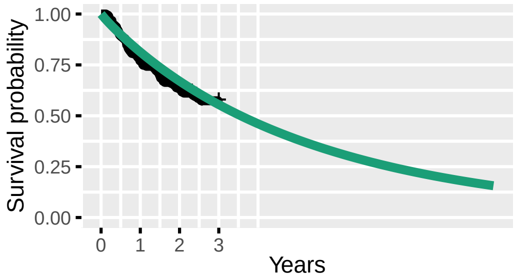

Hazard function $h(t|\lambda, \alpha) = \lambda \alpha t^{\alpha - 1}$

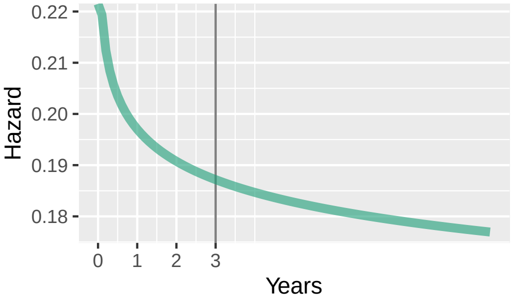

Parametric survival models: examples

Log-logistic distribution

Survivor function
$S(t|\lambda, \alpha) = 1/(1 + (t/b)^a)$

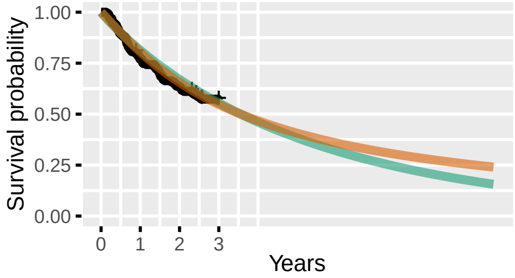

Hazard function $h(t|\lambda, \alpha) =$

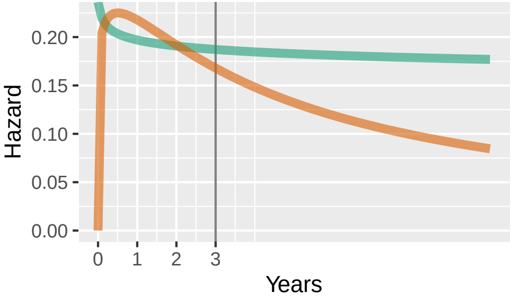

Parametric survival models: examples

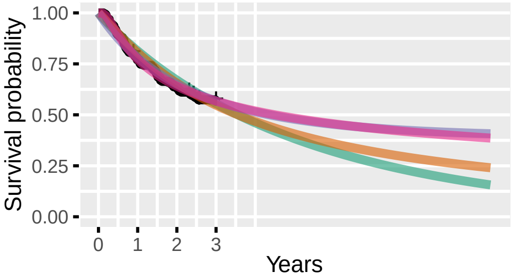

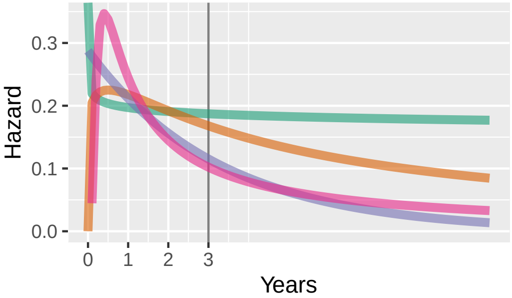

Parametric survival models: examples

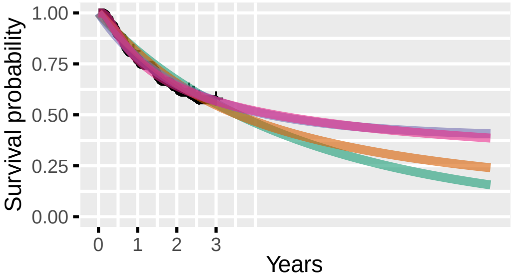

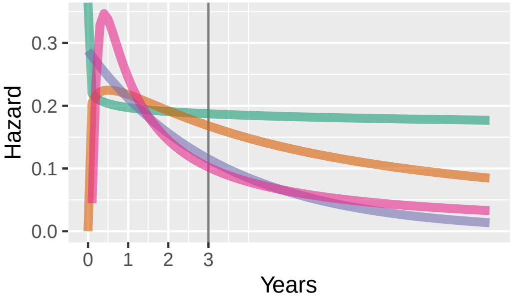

Fit a set of models and judge which :

- fits the data best in the short term? (easy, e.g. AIC)
- gives more plausible risk changes in the long term? harder

External information about longer-term survival

Observed data

- survival on "standard of care" arm same as recent patients in a disease registry?
- survival of people with the disease cannot be better than comparable people in the general population?

Clinical judgements about the mechanism

- e.g. some people get cured of the disease, so disease-specific hazard of death reduces to zero

Formally-elicited judgements

e.g. probability of survival to 10 years, given survival to 5 years

Use such information in a transparent, statistically-principled way

survextrap R package

https://chjackson.github.io/survextrap

Jackson, BMC Med Res Meth, 2023

Design goals:

1. Incorporate all available data: *Bayesian evidence synthesis*
2. Fit the data as well as possible: *Spline modelling*
3. Quantify uncertainty: *Bayesian*
4. Be easy to use: *R package wraps Stan, single commands to fit model / extract outputs*
5. Make key assumptions transparent: *user states these in their command, computer does the hard work*

Rapid tour. . .

survextrap R package

https://chjackson.github.io/survextrap

Jackson, BMC Med Res Meth, 2023

Design goals:

1. Incorporate all available data: *Bayesian evidence synthesis*
2. Fit the data as well as possible: *Spline modelling*
3. Quantify uncertainty: *Bayesian*
4. Be easy to use: *R package wraps Stan, single commands to fit model / extract outputs*
5. Make key assumptions transparent: *user states these in their command, computer does the hard work*

Rapid tour. . .

survextrap R package

https://chjackson.github.io/survextrap

Jackson, BMC Med Res Meth, 2023

Design goals:

1. Incorporate all available data: *Bayesian evidence synthesis*
2. Fit the data as well as possible: *Spline modelling*
3. Quantify uncertainty: *Bayesian*
4. Be easy to use: *R package wraps Stan, single commands to fit model / extract outputs*
5. Make key assumptions transparent: *user states these in their command, computer does the hard work*

Rapid tour. . .

Spline-based, flexible model for the hazard

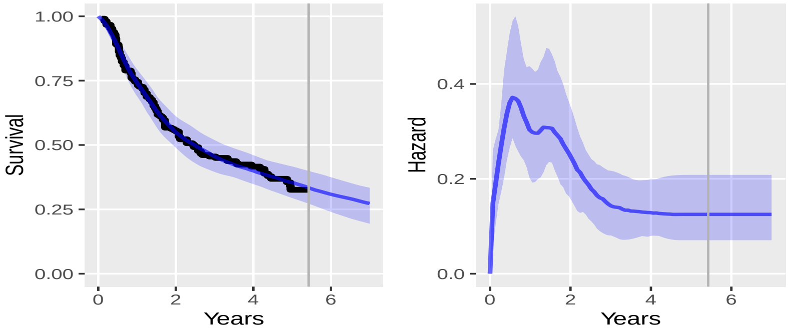

Running example: head and neck cancer trial data from Guyot et al. 2017

Default model: M-spline

10 parameters, penalized to control overfitting

Flexible hazard within data, constant beyond data

```
mod1 <- survextrap(Surv(years, d) ~ 1, data=dat)
```

Spline-based, flexible model for the hazard

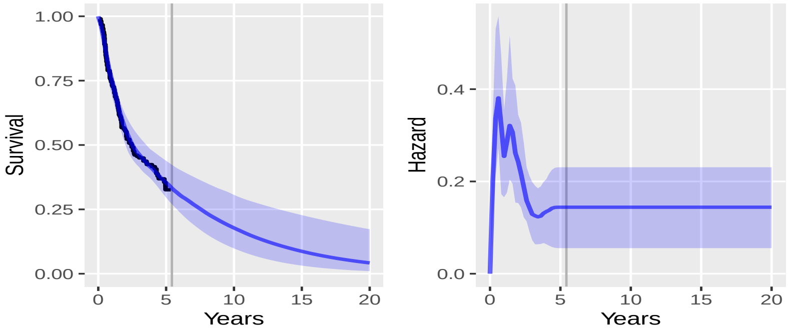

Spline-based, flexible model for the hazard

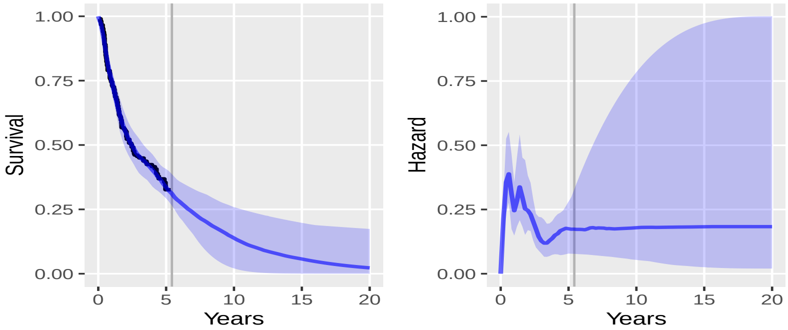

Allow hazard to vary to 20 years, by adding a spline "knot".

Weak assumption: hazard is smooth — short-term extrapolation influenced weakly by latest data.

Posterior represents "structural" uncertainty due to lack of data

```
mod2 <- survextrap(Surv(years, d) ~ 1, add_knots=20)
```

Including data external to the trial

Supplied as aggregate conditional survival counts over any number of different time intervals. Example:

Follow-up period		Number		Covariates
From t	To u	Alive at t	Still alive at u	
5	10	358	325	...
10	15	308	285	...
etc.				

- Registry data, population data, and elicited judgements... can often be expressed in this form
- Each count of survivors is a binomal outcome, with probability defined by the spline model and covariates
- Bayesian evidence synthesis: posterior for parameters determined given individual and external data together

Including data external to the trial

Supplied as aggregate conditional survival counts over any number of different time intervals. Example:

Follow-up period		Number		Covariates
From t	To u	Alive at t	Still alive at u	
5	10	358	325	...
10	15	308	285	...
etc.				

- Registry data, population data, and elicited judgements... can often be expressed in this form
- Each count of survivors is a binomal outcome, with probability defined by the spline model and covariates
- Bayesian evidence synthesis: posterior for parameters determined given individual and external data together

Survival modelling with external data: example

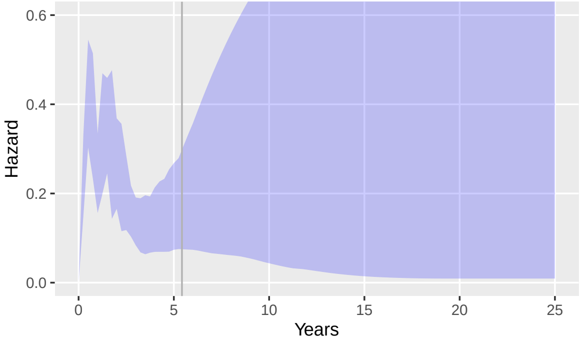

With no external data, and just a smoothness assumption, estimated long-term hazard is extremely uncertain

```
mod_seer <- survextrap(Surv(years, d) ~ 1, data=dat,
                       add_knots=20)
```

Survival modelling with external data: example

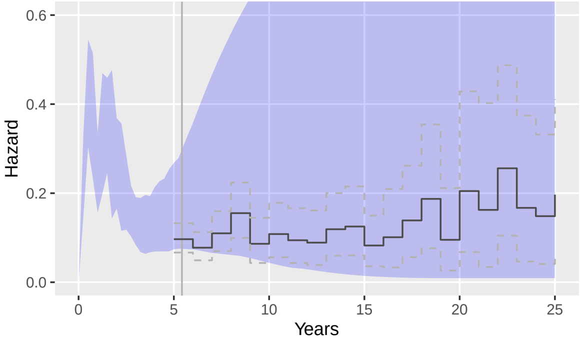

Registry data give annual survival rates from 5 to 25 years.

Survival modelling with external data: example

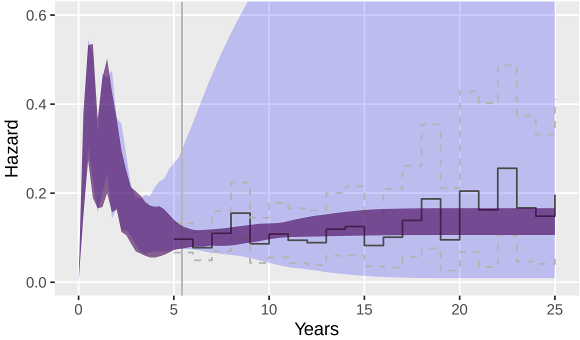

Registry data give annual survival rates from 5 to 25 years.

- Enables a confident "extrapolation" of the hazard for control group survival

```
mod_seer <- survextrap(Surv(years, d) ~ 1, data=dat,
                       external=ext_dat, add_knots=20)
```

Survival modelling with external data: example

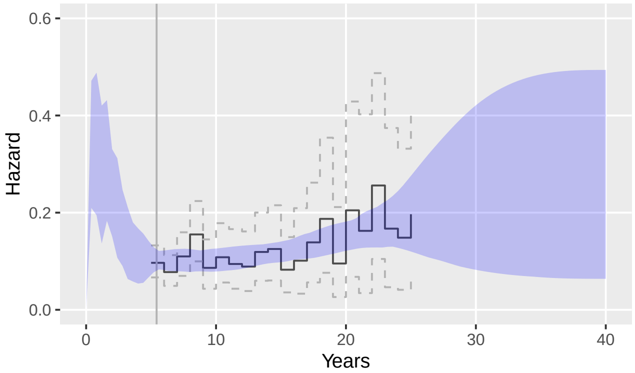

Extrapolating up to 40 years now...

```
mod_seer <- survextrap(Surv(years, d) ~ 1, data=dat,
                       external=ext_dat,
                       add_knots=c(20,30,40))
```

Additive hazards model

Can supply a known background hazard $h_b(t)$

- e.g. national mortality statistics for people of comparable age/sex

Hazard for trial patients modelled as additive
Flexible spline model placed on the excess hazard $h_e(t)$ instead of $h(t)$

$$h(t) = h_b(t) + h_e(t)$$

i.e. no lower than population hazard

Additive hazards model

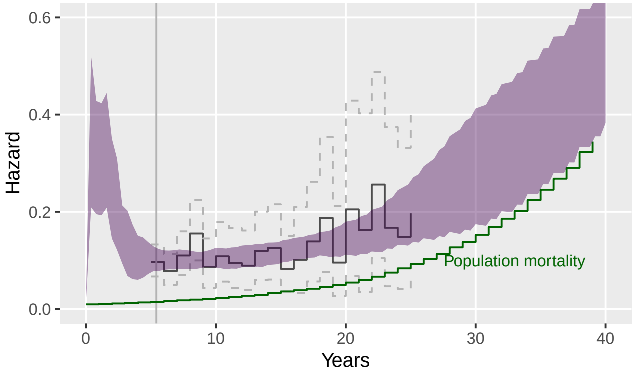


```
mod_seer <- survextrap(Surv(years, d) ~ 1, data=dat,
                       external=ext_dat, bh=back_dat,
                       add_knots=c(20,30,40))
```

Elicited survival probabilities with credible intervals can be interpreted as aggregate survival counts

- Example: suppose we judge that over some interval, the survival probability is 0.72 (0.69, 0.75)
- Find $Beta(a, b)$ distribution that best fits this belief

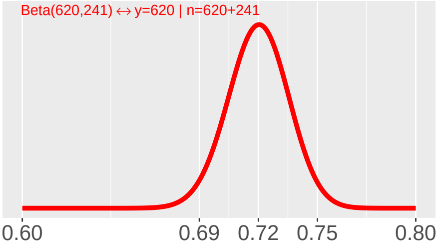

- Find Binomial data $(y|n)$ that would give the (conjugate) posterior $Beta(a, b)$ when combined with a vague prior

Including expert judgements on long-term survival

Elicited survival probabilities with credible intervals can be interpreted as aggregate survival counts

- Example: suppose we judge that over some interval, the survival probability is 0.72 (0.69, 0.75)
- Find $Beta(a, b)$ distribution that best fits this belief

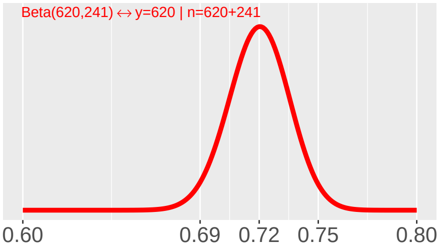

- Find Binomial data $(y|n)$ that would give the (conjugate) posterior $Beta(a, b)$ when combined with a vague prior
- Treat $(y|n)$ as "external data" to be analysed jointly with the trial data

Including expert judgements on long-term survival

Elicited survival probabilities with credible intervals can be interpreted as aggregate survival counts

- Example: suppose we judge that over some interval, the survival probability is 0.72 (0.69, 0.75)
- Find $Beta(a, b)$ distribution that best fits this belief

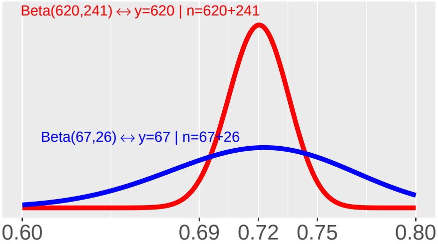

- Find Binomial data $(y|n)$ that would give the (conjugate) posterior $Beta(a, b)$ when combined with a vague prior
- Treat $(y|n)$ as "external data" to be analysed jointly with the trial data

Including expert judgements on long-term survival

Elicited survival probabilities with credible intervals can be interpreted as aggregate survival counts

- Example: suppose we judge that over some interval, the survival probability is 0.72 (0.69, 0.75)
- Find $Beta(a, b)$ distribution that best fits this belief

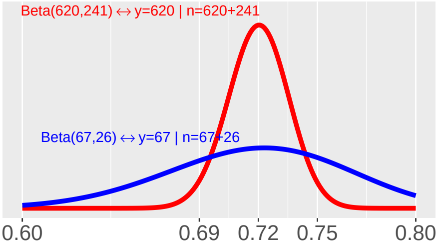

  - Find Binomial data $(y|n)$ that would give the (conjugate) posterior $Beta(a, b)$ when combined with a vague prior
  - Treat $(y|n)$ as "external data" to be analysed jointly with the trial data

Model combining trial data and expert judgement

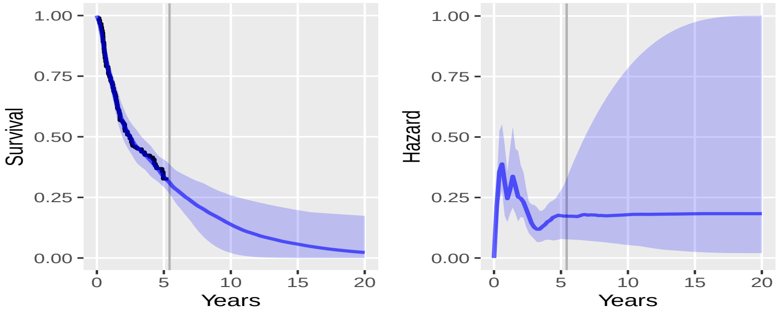

Trial data only

```
mod_seer <- survextrap(Surv(years, d) ~ 1, data=dat,
                       add_knots=c(20))
```

Model combining trial data and expert judgement

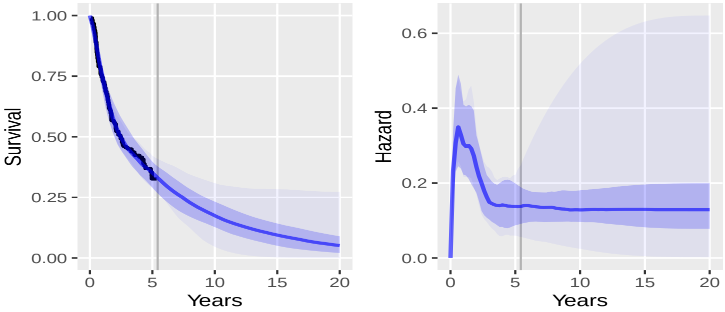

Trial data + elicited probability that patients who survive the trial (5 years) are still alive at 20 years: 12% (95% CI 4% to 25%).

```
mod_seer <- survextrap(Surv(years, d) ~ 1, data=dat,
  external = data.frame(r=5, n=35, start=5, stop=20),
  add_knots=c(20))
```

Treatment or covariate effects

Proportional hazards model

```
mod_seer <- survextrap(Surv(years, d) ~ treat + age,
                      data=dat, external=ext_dat)
```

Flexible non-proportional hazards model

```
mod_seer <- survextrap(Surv(years, d) ~ treat + age,
                      data=dat, external=ext_dat,
                      nonprop=TRUE)
```

Additionally can impose a waning effect when predicting

- log hazard ratio for a treatment decreases from modelled value to zero over some time period
- for sensitivity analysis to what happens after trial data

Treatment or covariate effects

Proportional hazards model

```
mod_seer <- survextrap(Surv(years, d) ~ treat + age,
                       data=dat, external=ext_dat)
```

Flexible non-proportional hazards model

```
mod_seer <- survextrap(Surv(years, d) ~ treat + age,
                       data=dat, external=ext_dat,
                       nonprop=TRUE)
```

Additionally can impose a waning effect when predicting

- log hazard ratio for a treatment decreases from modelled value to zero over some time period
- for sensitivity analysis to what happens after trial data

Treatment or covariate effects

Proportional hazards model

```
mod_seer <- survextrap(Surv(years, d) ~ treat + age,
                       data=dat, external=ext_dat)
```

Flexible non-proportional hazards model

```
mod_seer <- survextrap(Surv(years, d) ~ treat + age,
                       data=dat, external=ext_dat,
                       nonprop=TRUE)
```

Additionally can impose a waning effect when predicting

- log hazard ratio for a treatment decreases from modelled value to zero over some time period
- for sensitivity analysis to what happens after trial data

Further information

Technical details, nuisance choices

- Prior distributions, detailed specification of spline (knots, smoothness penalty...), can all be customised
- Sensible package defaults demonstrated by simulation

Detailed documentation and worked examples

https://chjackson.github.io/survextrap

Jackson, BMC Med Res Meth, 2023

Simulation studies

Timmins I., Torabi F., Jackson C., Lambert P., Sweeting M. (2025) *Simulation-based assessment of a Bayesian survival model with flexible baseline hazard and time-dependent effects.* 10.48550/arXiv.2503.21388

and ongoing work

Further information

Technical details, nuisance choices

- Prior distributions, detailed specification of spline (knots, smoothness penalty...), can all be customised
- Sensible package defaults demonstrated by simulation

Detailed documentation and worked examples

https://chjackson.github.io/survextrap

Jackson, BMC Med Res Meth, 2023

Simulation studies

Timmins I., Torabi F., Jackson C., Lambert P., Sweeting M. (2025) *Simulation-based assessment of a Bayesian survival model with flexible baseline hazard and time-dependent effects.* 10.48550/arXiv.2503.21388

and ongoing work