System Reliability Estimation under Prior-Data Conflict

Gero Walter[1], Frank Coolen[2], Simme Douwe Flapper[1]

[1]Eindhoven University of Technology, Eindhoven, NL
[2]Durham University, Durham, UK

g.m.walter@tue.nl

Technische Universiteit Eindhoven University of Technology

Durham University

DINALOG Dutch Institute for Advanced Logistics

2015-11-03

Technische Universiteit Eindhoven University of Technology

Setting: a prototype system

We want to learn about the system reliability $R_{\text{sys}}(t) = P(T_{\text{sys}} > t)$ based on

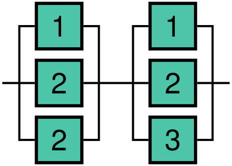

Technische Universiteit Eindhoven University of Technology

Setting: a prototype system

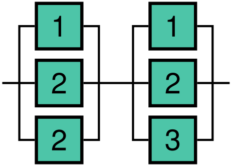

We want to learn about the system reliability $R_{\text{sys}}(t) = P(T_{\text{sys}} > t)$ based on

- system run until time t_{now}:

  ℓ observations, each being either
  a failure time t_j or a censoring time $t_j^+ = t_{\text{now}}$

Technische Universiteit Eindhoven University of Technology

Setting: a prototype system

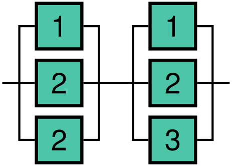

We want to learn about the system reliability $R_{\text{sys}}(t) = P(T_{\text{sys}} > t)$ based on

- system run until time t_{now}:

  ℓ observations, each being either
  a failure time t_j or a censoring time $t_j^+ = t_{\text{now}}$

- cautious assumptions on component reliability:

  expert information,
  e.g. from the component manufacturers
  which we don't trust entirely

TU/e Technische Universiteit Eindhoven University of Technology

Setting: a prototype system

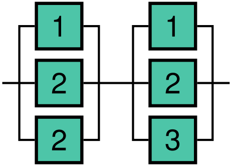

We want to learn about the system reliability $R_{\text{sys}}(t) = P(T_{\text{sys}} > t)$ based on

- system run until time t_{now}:

  ℓ observations, each being either
  a failure time t_j or a censoring time $t_j^+ = t_{\text{now}}$
- cautious assumptions on component reliability:

  expert information,
  e.g. from the component manufacturers
  which we don't trust entirely

How to combine these two information sources?

Bayesian inference

expert info + data $\rightarrow$ complete picture

Bayesian inference

expert info	+	data	$\rightarrow$	complete picture	
prior distribution	+	likelihood	$\rightarrow$	posterior distribution	
$p(\lambda)$	$\times$	$p_c(\boldsymbol{t} \mid \lambda)$	$\propto$	$p(\lambda \mid \boldsymbol{t})$	► Bayes' Rule

Bayesian inference

expert info	+	data	$\rightarrow$	complete picture	
prior distribution	+	likelihood	$\rightarrow$	posterior distribution	
$p(\lambda)$	$\times$	$p_c(\boldsymbol{t} \mid \lambda)$	$\propto$	$p(\lambda \mid \boldsymbol{t})$	► Bayes' Rule
$\downarrow$		$\downarrow$		$\downarrow$	
inverse Gamma prior		Weibull with fixed shape κ		inverse Gamma posterior	► conjugacy
$\lambda \sim \mathrm{IG}(\alpha^{(0)}, \beta^{(0)})$		$\boldsymbol{t} \mid \lambda \sim \mathrm{Wei}_\kappa(\lambda)$		$\lambda \mid \boldsymbol{t} \sim \mathrm{IG}(\alpha^{(n)}, \beta^{(n)})$	

Bayesian inference

expert info	+	data	$\rightarrow$	complete picture	
prior distribution	+	likelihood	$\rightarrow$	posterior distribution	
$p(\lambda)$	$\times$	$p_c(\boldsymbol{t} \mid \lambda)$	$\propto$	$p(\lambda \mid \boldsymbol{t})$	▶ Bayes' Rule
$\downarrow$		$\downarrow$		$\downarrow$	
inverse Gamma prior		Weibull with fixed shape κ		inverse Gamma posterior	▶ conjugacy
$\lambda \sim \mathrm{IG}(\alpha^{(0)}, \beta^{(0)})$		$\boldsymbol{t} \mid \lambda \sim \mathrm{Wei}_\kappa(\lambda)$		$\lambda \mid \boldsymbol{t} \sim \mathrm{IG}(\alpha^{(n)}, \beta^{(n)})$	

- makes learning about component reliability tractable, just update parameters: $\alpha^{(0)} \rightarrow \alpha^{(n)}$, $\beta^{(0)} \rightarrow \beta^{(n)}$
- conjugacy holds also for censored observations
- closed form for system reliability function $R_{\mathsf{sys}}(t \mid \boldsymbol{t})$

Prior-data conflict

What if expert information and data tell different stories?

Prior-data conflict

What if expert information and data tell different stories?

- reparametrization helps to understand effect of prior-data conflict:

$$n^{(0)} = \alpha^{(0)} - 1\,, \qquad y^{(0)} = \beta^{(0)}/(\alpha^{(0)} - 1)\,, \qquad \text{where}$$

$$n^{(n)} = n^{(0)} + n\,, \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n}\, y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{1}{n} \textstyle\sum_{j=1}^{n} t_j^{\kappa}$$

Prior-data conflict

What if expert information and data tell different stories?

- reparametrization helps to understand effect of prior-data conflict:

$$n^{(0)} = \alpha^{(0)} - 1\,, \qquad y^{(0)} = \beta^{(0)}/(\alpha^{(0)} - 1)\,, \qquad \text{where}$$

$$n^{(n)} = n^{(0)} + n\,, \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n}\, y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{1}{n} \sum_{j=1}^{n} t_j^{\kappa}$$

$$\boxed{y^{(0)} = \mathrm{E}[\lambda]}$$

Prior-data conflict

What if expert information and data tell different stories?

- reparametrization helps to understand effect of prior-data conflict:

$$n^{(0)} = \alpha^{(0)} - 1\,, \qquad y^{(0)} = \beta^{(0)}/(\alpha^{(0)} - 1)\,, \qquad \text{where}$$

$$n^{(n)} = n^{(0)} + n\,, \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n}\, y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{1}{n} \textstyle\sum_{j=1}^{n} t_j^{\kappa}$$

$$y^{(0)} = \mathrm{E}[\lambda] \qquad y^{(n)} = \mathrm{E}[\lambda \mid \boldsymbol{t}]$$

Prior-data conflict

What if expert information and data tell different stories?

- reparametrization helps to understand effect of prior-data conflict:

$$n^{(0)} = \alpha^{(0)} - 1\,, \qquad y^{(0)} = \beta^{(0)}/(\alpha^{(0)} - 1)\,, \qquad \text{where}$$

$$n^{(n)} = n^{(0)} + n\,, \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n}\, y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \underbrace{\frac{1}{n} \textstyle\sum_{j=1}^{n} t_j^{\kappa}}$$

$y^{(0)} = \mathrm{E}[\lambda]$ | $y^{(n)} = \mathrm{E}[\lambda \mid \boldsymbol{t}]$ | ML estimator $\hat{\lambda}$

Prior-data conflict

What if expert information and data tell different stories?

- reparametrization helps to understand effect of prior-data conflict:

$$n^{(0)} = \alpha^{(0)} - 1\,, \qquad y^{(0)} = \beta^{(0)}/(\alpha^{(0)} - 1)\,, \qquad \text{where}$$

$$n^{(n)} = n^{(0)} + n\,, \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n}\, y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{1}{n} \sum_{j=1}^{n} t_j^{\kappa}$$

$n^{(0)}$ = pseudocounts

$y^{(0)} = \mathrm{E}[\lambda]$

$y^{(n)} = \mathrm{E}[\lambda \mid \boldsymbol{t}]$

ML estimator $\hat{\lambda}$

Prior-data conflict

What if expert information and data tell different stories?

- reparametrization helps to understand effect of prior-data conflict:

$$n^{(0)} = \alpha^{(0)} - 1\,, \qquad y^{(0)} = \beta^{(0)}/(\alpha^{(0)} - 1)\,, \qquad \text{where}$$

$$n^{(n)} = n^{(0)} + n\,, \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n}\, y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{1}{n} \sum_{j=1}^{n} t_j^{\kappa}$$

$n^{(0)}$ = pseudocounts

$y^{(0)} = \mathrm{E}[\lambda]$

$y^{(n)} = \mathrm{E}[\lambda \mid \boldsymbol{t}]$

ML estimator $\hat{\lambda}$

$\mathrm{E}[\lambda \mid \boldsymbol{t}]$ is a weighted average of $\mathrm{E}[\lambda]$ and $\hat{\lambda}$!

Prior-data conflict example

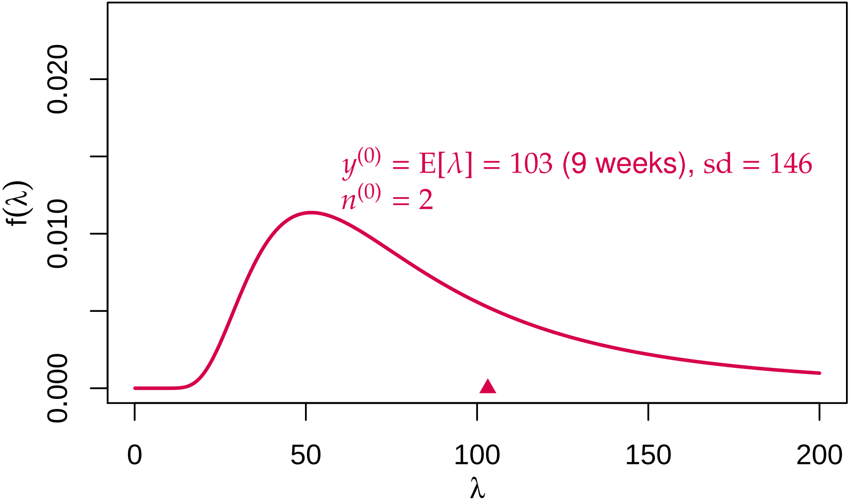

Technische Universiteit **Eindhoven** University of Technology

Prior-data conflict example

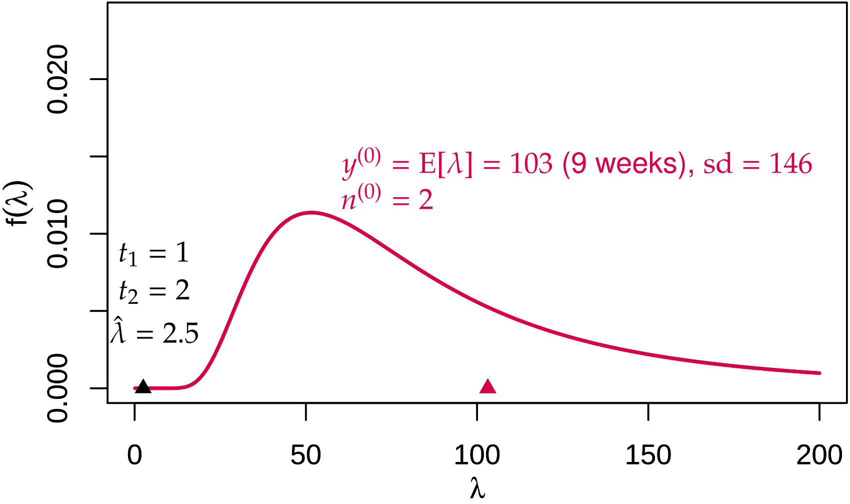

Prior-data conflict example

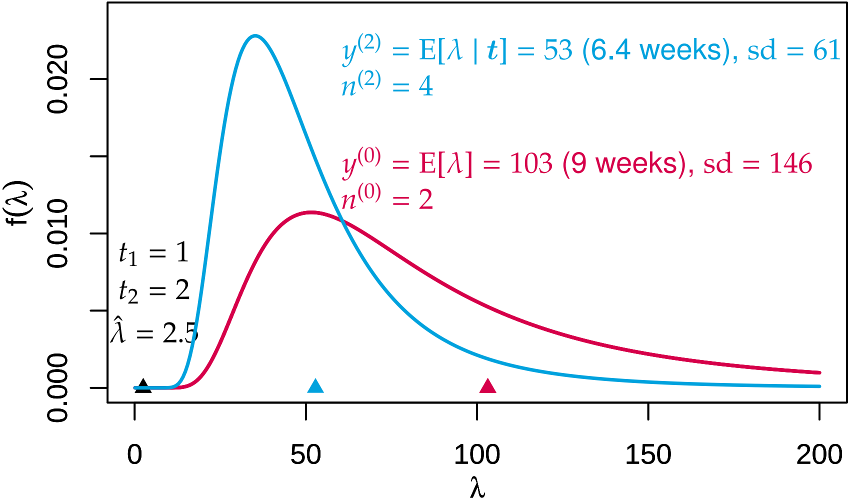

Technische Universiteit Eindhoven University of Technology

Prior-data conflict example

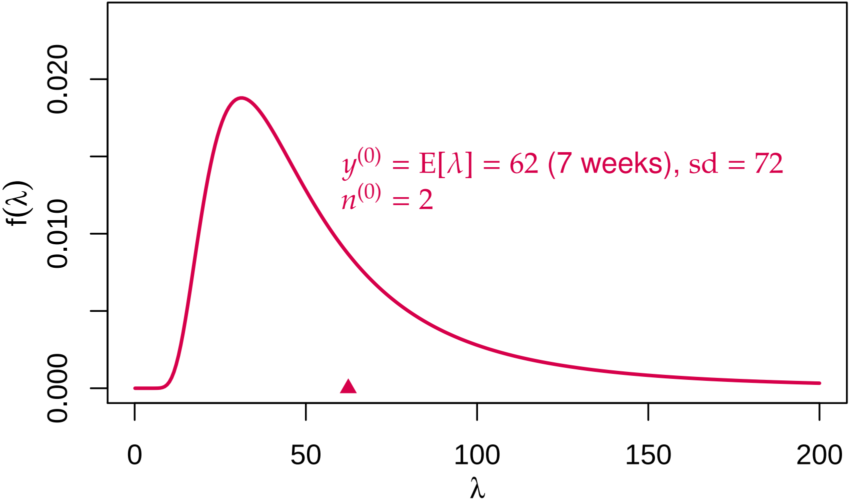

Prior-data conflict example

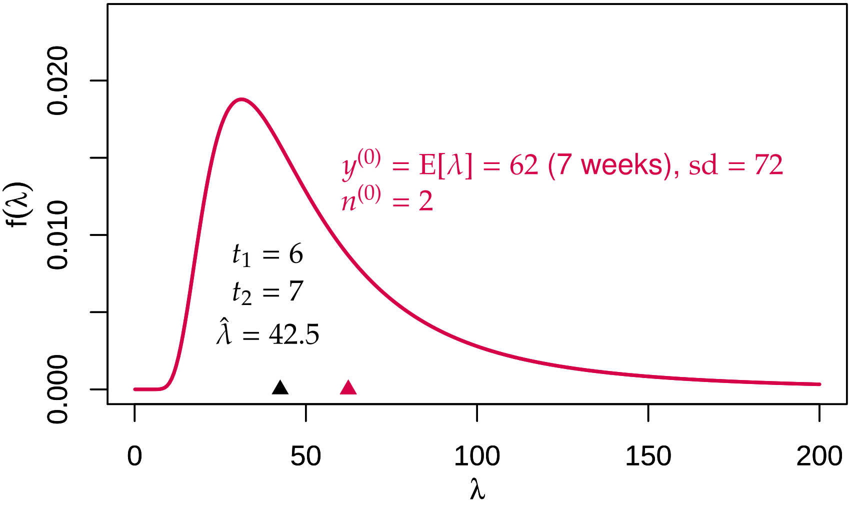

Prior-data conflict example

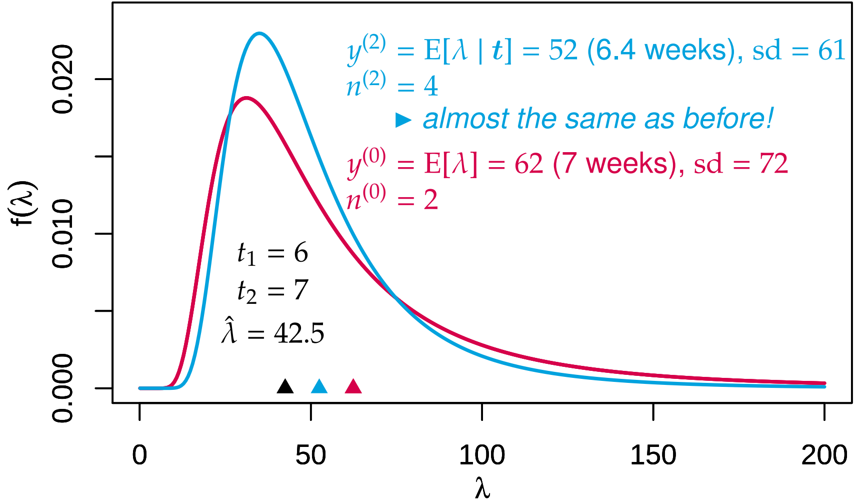

Imprecise / interval probability

- Averaging property holds *for all conjugate models* (!)
  Can we mitigate this and still keep tractability?

Imprecise / interval probability

- Averaging property holds *for all conjugate models* (!)
  Can we mitigate this and still keep tractability?
- Reliability function $R(t)$ is a collection of probability statements:
  $R(t) =$ probability that the system survives past t.
  How can we express uncertainty
  about these probability statements?

Imprecise / interval probability

- Averaging property holds *for all conjugate models* (!)
  Can we mitigate this and still keep tractability?
- Reliability function $R(t)$ is a collection of probability statements:
  $R(t) =$ probability that the system survives past t.
  How can we express uncertainty
  about these probability statements?
- **Add imprecision as new modelling dimension:**
  Sets of priors model uncertainty in probability statements and allow to better model partial or vague information on λ.

TU/e Technische Universiteit Eindhoven University of Technology

Imprecise / interval probability

- Averaging property holds *for all conjugate models* (!)
  Can we mitigate this and still keep tractability?
- Reliability function $R(t)$ is a collection of probability statements:
  $R(t) =$ probability that the system survives past t.
  How can we express uncertainty
  about these probability statements?
- **Add imprecision as new modelling dimension:**
  Sets of priors model uncertainty in probability statements and allow to better model partial or vague information on λ.
- Separate uncertainty *whithin the model* (reliability statements) from uncertainty *about the model* (which parameters).

Imprecise / interval probability

- Averaging property holds *for all conjugate models* (!)
  Can we mitigate this and still keep tractability?
- Reliability function $R(t)$ is a collection of probability statements:
  $R(t) =$ probability that the system survives past t.
  How can we express uncertainty
  about these probability statements?
- **Add imprecision as new modelling dimension:**
  Sets of priors model uncertainty in probability statements and allow to better model partial or vague information on λ.
- Separate uncertainty *whithin the model* (reliability statements) from uncertainty *about the model* (which parameters).
- Can also be seen as systematic sensitivity analysis or robust Bayesian approach.

Sets of prior distributions

Uncertainty about probability statements

smaller sets = more precise probability statements

Lottery A

Number of winning tickets:
exactly known as 5 out of 100

→ $P(\text{win}) = 5/100$

Lottery B

Number of winning tickets:
not exactly known, supposedly
between 1 and 7 out of 100

→ $P(\text{win}) = [1/100, 7/100]$

Sets of prior distributions

Uncertainty about probability statements

smaller sets = more precise probability statements

Lottery A
Number of winning tickets:
exactly known as 5 out of 100
→ $P(\text{win}) = 5/100$

Lottery B
Number of winning tickets:
not exactly known, supposedly
between 1 and 7 out of 100
→ $P(\text{win}) = [1/100, 7/100]$

Let parameters $(n^{(0)}, y^{(0)})$ vary in a set $\Pi^{(0)}$ → set of priors

Sets of priors $\rightarrow$ sets of posteriors by updating element by element:
GBR (Walley 1991) ensures *coherence* (a consistency property)

Sets of prior distributions

Uncertainty about probability statements

smaller sets = more precise probability statements

Lottery A
Number of winning tickets:
exactly known as 5 out of 100
➡ $P(\text{win}) = 5/100$

Lottery B
Number of winning tickets:
not exactly known, supposedly
between 1 and 7 out of 100
➡ $P(\text{win}) = [1/100, 7/100]$

Let parameters $(n^{(0)}, y^{(0)})$ vary in a set $\Pi^{(0)}$ ➡ set of priors

Sets of priors → sets of posteriors by updating element by element:
GBR (Walley 1991) ensures *coherence* (a consistency property)

Walter and Augustin (2009), Walter (2013):
$\Pi^{(0)} = [\underline{n}^{(0)}, \overline{n}^{(0)}] \times [\underline{y}^{(0)}, \overline{y}^{(0)}]$
gives tractability & meaningful reaction to prior-data conflict:

- larger set of posteriors
- more imprecise / cautious probability statements

Technische Universiteit Eindhoven University of Technology

/

Sets of prior distributions: examples

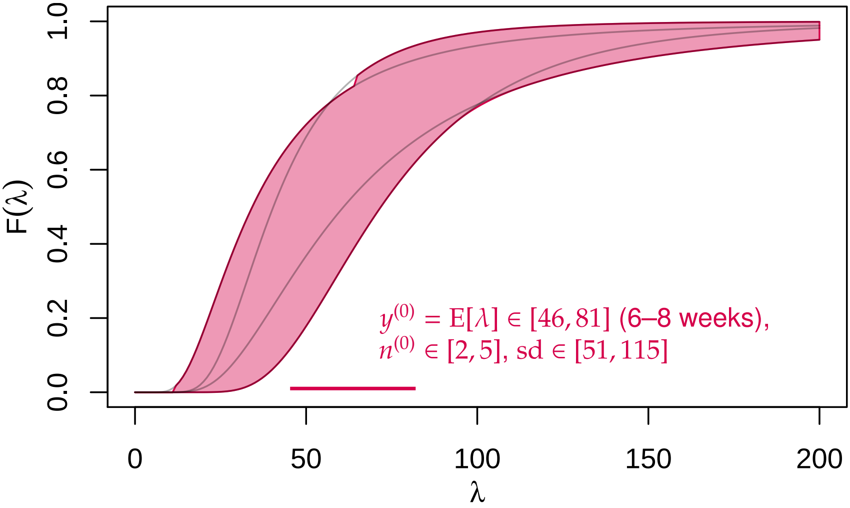

Sets of prior distributions: examples

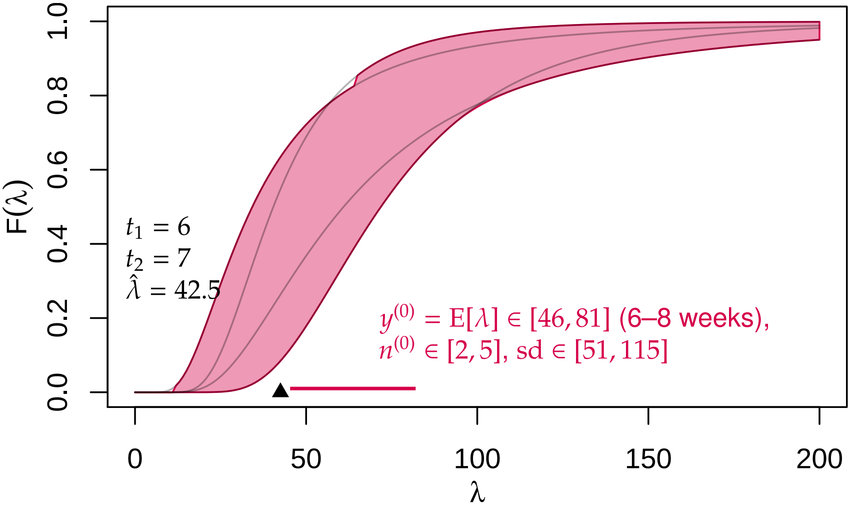

Sets of prior distributions: examples

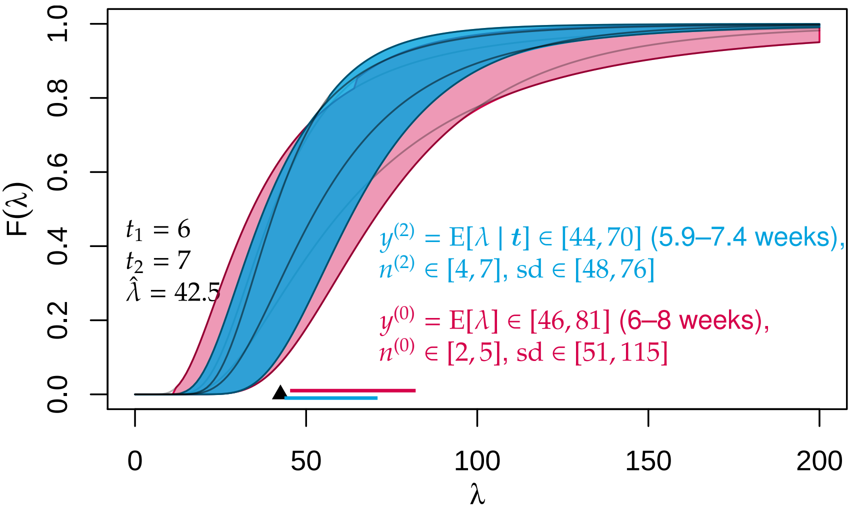

Sets of prior distributions: examples

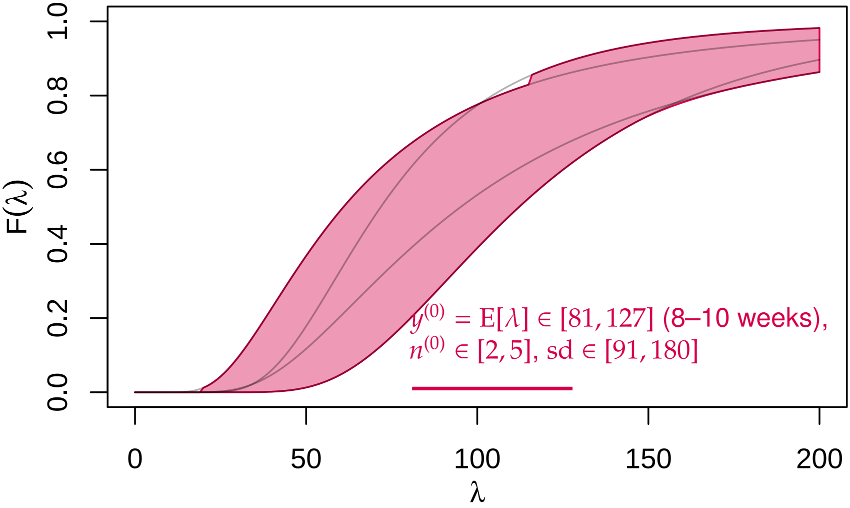

TU/e Technische Universiteit Eindhoven University of Technology

Sets of prior distributions: examples

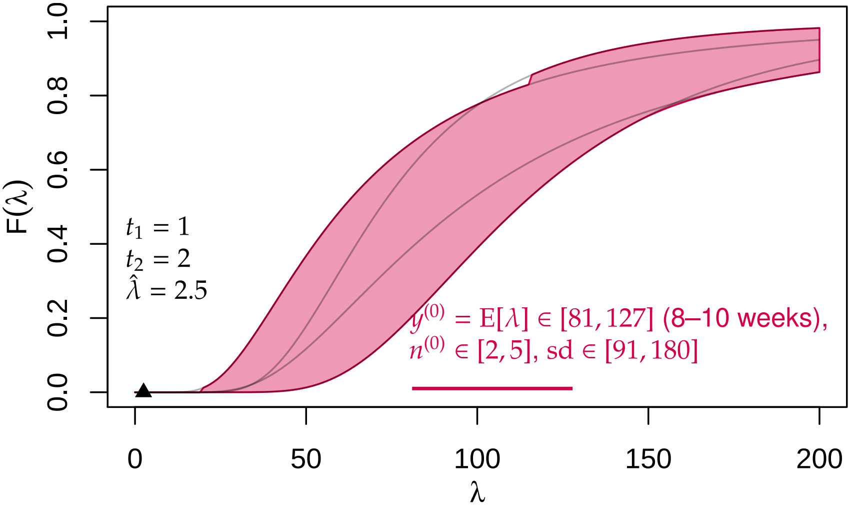

TU/e Technische Universiteit Eindhoven University of Technology

Sets of prior distributions: examples

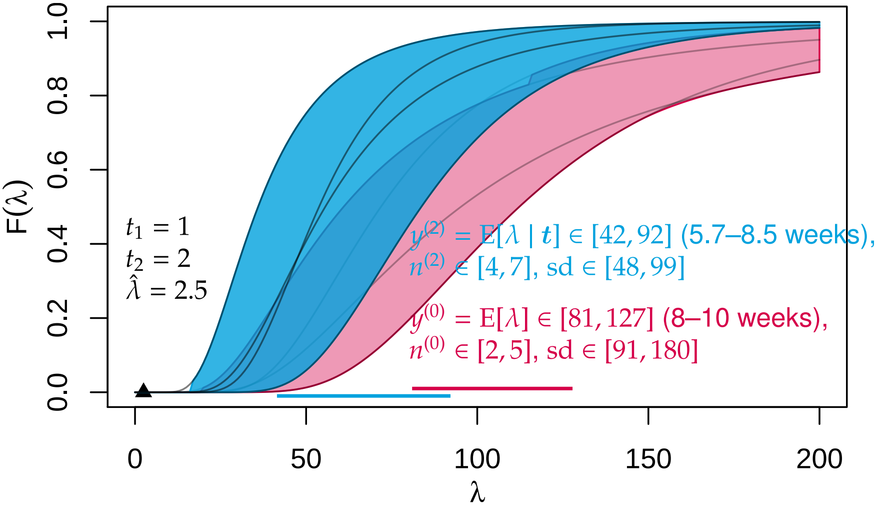

Sets of prior distributions: examples

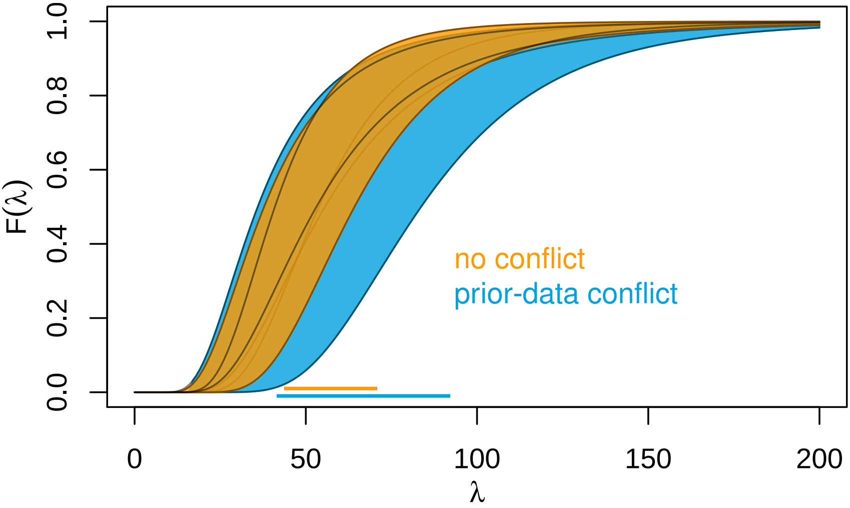

System reliability

- Closed form for the system reliability via the survival signature:

$$P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$
$$= \sum_{l_1=0}^{n_1-e_1} \cdots \sum_{l_K=0}^{n_K-e_K} \Phi(l_1, \ldots, l_K) \prod_{k=1}^{K} P(C_t^k = l_k \mid n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k)$$

System reliability

- Closed form for the system reliability via the survival signature:

$$P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

$$= \sum_{l_1=0}^{n_1-e_1} \cdots \sum_{l_K=0}^{n_K-e_K} \Phi(l_1,\ldots,l_K) \prod_{k=1}^{K} P(C_t^k = l_k \mid n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k)$$

Survival signature $\Phi(l_1,\ldots,l_K)$
(Coolen and Coolen-Maturi 2012)
$= P(\text{system functions} \mid \{l_k$ k 's function$\}^{1:K})$

l_1	l_2	l_3	Φ
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	0.67
0	2	0	0.67

l_1	l_2	l_3	Φ
0	2	1	1
0	3	0	1
0	3	1	1
1	0	0	0
⋮	⋮	⋮	⋮

System reliability

- Closed form for the system reliability via the survival signature:

$$P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$
$$= \sum_{l_1=0}^{n_1-e_1} \cdots \sum_{l_K=0}^{n_K-e_K} \Phi(l_1,\ldots,l_K) \prod_{k=1}^{K} P(C_t^k = l_k \mid n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k)$$

Survival signature $\Phi(l_1,\ldots,l_K)$
(Coolen and Coolen-Maturi 2012)
$= P$(system functions | $\{l_k$ k 's function$\}^{1:K}$)

l_1	l_2	l_3	Φ
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	0.67
0	2	0	0.67

l_1	l_2	l_3	Φ
0	2	1	1
0	3	0	1
0	3	1	1
1	0	0	0
⋮	⋮	⋮	⋮

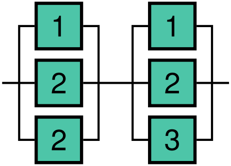

System reliability

- Closed form for the system reliability via the survival signature:

$$P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right) = \sum_{l_1=0}^{n_1-e_1} \cdots \sum_{l_K=0}^{n_K-e_K} \Phi(l_1, \ldots, l_K) \prod_{k=1}^{K} P(C_t^k = l_k \mid n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k)$$

Survival signature $\Phi(l_1, \ldots, l_K)$
(Coolen and Coolen-Maturi 2012)
$= P(\text{system functions} \mid \{l_k$ k's function$\}^{1:K})$

l_1	l_2	l_3	Φ
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	0.67
0	2	0	0.67
0	2	1	1
0	3	0	1
0	3	1	1
1	0	0	0
⋮	⋮	⋮	⋮

Posterior predictive probability that l_k of the $n_k - e_k$ surviving k's function at time t:

$$\binom{n_k-e_k}{l_k} \int [P(t > T \mid T > t_{\text{now}}, \lambda_k)]^{l_k} [P(t \leq T \mid T > t_{\text{now}}, \lambda_k)]^{n_k-e_k-l_k} f_{\lambda_k|\ldots}(\lambda_k \mid n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k)\, d\theta$$

(integral can be solved analytically)

Technische Universiteit Eindhoven University of Technology

System reliability: bounds

- Lower / upper bound through optimization for each t:

$$\underline{R}_{\mathsf{sys}}\left(t \mid \{\mathrm{I\!\Pi}_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right) = \min_{\mathrm{I\!\Pi}_1^{(0)},\ldots,\mathrm{I\!\Pi}_K^{(0)}} P\left(T_{\mathsf{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

System reliability: bounds

- Lower / upper bound through optimization for each t:

$$\underline{R}_{\mathsf{sys}}\left(t \mid \{\mathrm{I\!\Pi}_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right) = \min_{\mathrm{I\!\Pi}_1^{(0)},\dots,\mathrm{I\!\Pi}_K^{(0)}} P\left(T_{\mathsf{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

$$= \min_{n_1^{(0)},\dots,n_K^{(0)}} P\left(T_{\mathsf{sys}} > t \mid \{n_k^{(0)}, \underline{y}_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

System reliability: bounds

- Lower / upper bound through optimization for each t:

$$\underline{R}_{\text{sys}}\left(t \mid \{\Pi_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right) = \min_{\Pi_1^{(0)},\ldots,\Pi_K^{(0)}} P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

$$= \min_{n_1^{(0)},\ldots,n_K^{(0)}} P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, \underline{y}_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

$$\overline{R}_{\text{sys}}\left(t \mid \{\Pi_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right) = \max_{\Pi_1^{(0)},\ldots,\Pi_K^{(0)}} P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, y_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

$$= \max_{n_1^{(0)},\ldots,n_K^{(0)}} P\left(T_{\text{sys}} > t \mid \{n_k^{(0)}, \overline{y}_k^{(0)}, \boldsymbol{t}^k\}^{1:K}\right)$$

System reliability: examples

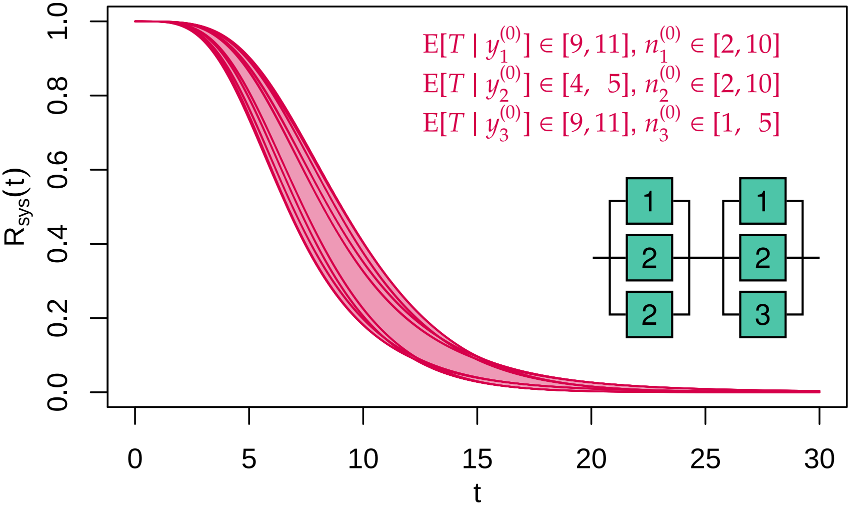

System reliability: examples

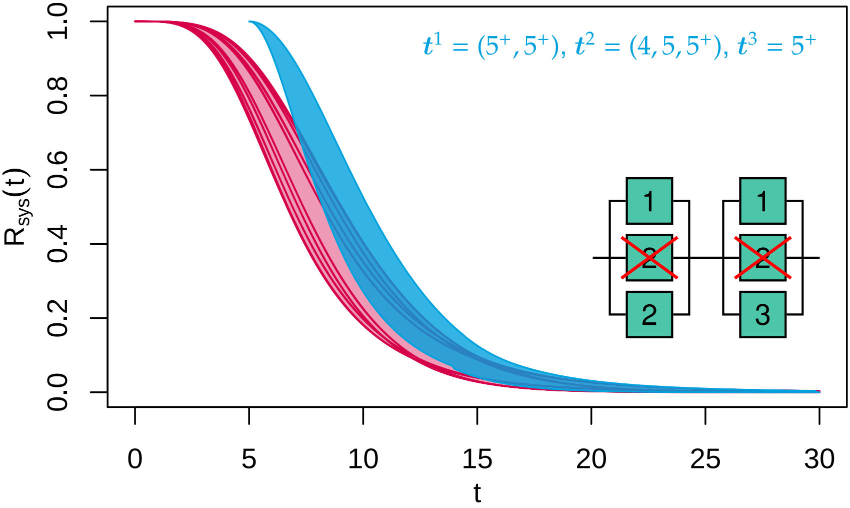

System reliability: examples

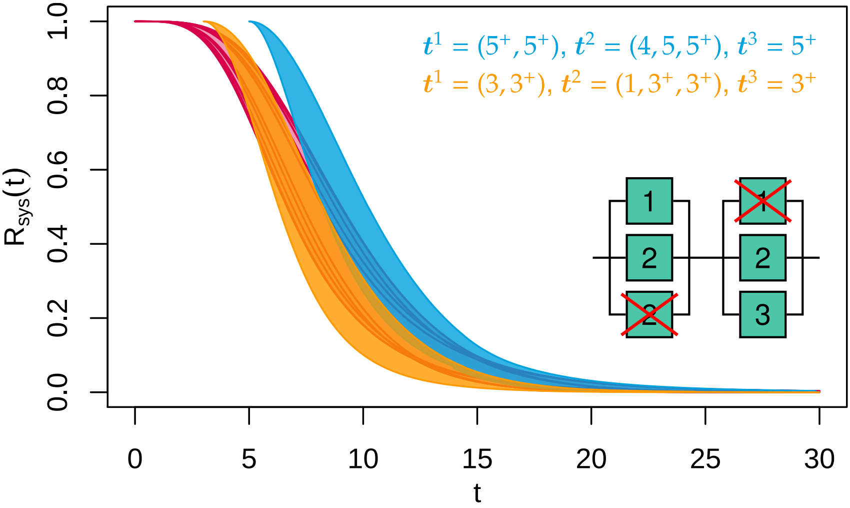

System reliability: examples

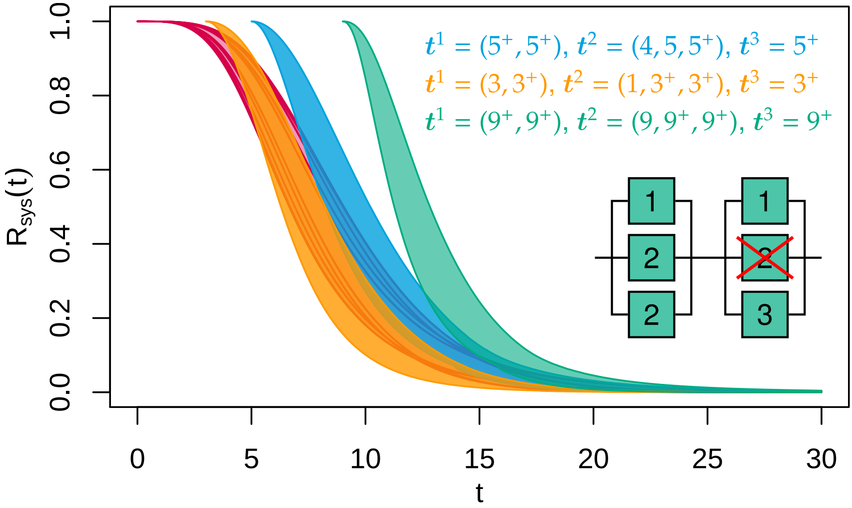

Summary & Outlook

Summary:

- Very limited data: Bayesian model with set of conjugate priors
- Set of system reliability functions reflects uncertainties from limited data (with censoring!) and vague expert information
- In particular, it reflects *prior-data conflict*

Summary & Outlook

Summary:

- Very limited data: Bayesian model with set of conjugate priors
- Set of system reliability functions reflects uncertainties from limited data (with censoring!) and vague expert information
- In particular, it reflects *prior-data conflict*

Next steps:

- Nonparametric model (drop Weibull assumption for component lifetimes)
- Allow dependence between components (common-cause failure, . . .)
- Use model for maintenance planning

References

Coolen, Frank P. A. and Tahani Coolen-Maturi (2012). “Generalizing the Signature to Systems with Multiple Types of Components”. In: *Complex Systems and Dependability*. Ed. by W. Zamojski et al. Vol. 170. Advances in Intelligent and Soft Computing. Springer, pp. 115–130. DOI: 10.1007/978-3-642-30662-4_8.

Walley, Peter (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.

Walter, Gero (2013). “Generalized Bayesian Inference under Prior-Data Conflict”. PhD thesis. Department of Statistics, LMU Munich. URL: http://edoc.ub.uni-muenchen.de/17059/.

Walter, Gero and Thomas Augustin (2009). “Imprecision and Prior-Data Conflict in Generalized Bayesian Inference”. In: *Journal of Statistical Theory and Practice* 3, pp. 255–271. DOI: 10.1080/15598608.2009.10411924.