Cautious Uncertainty Modelling in Common-Cause Failure Models with Sets of Conjugate Priors

Gero Walter

Durham University, UK

gero.walter@durham.ac.uk

24 November 2014

Outline

1. Common-cause failure modelling
   (joint work with Matthias Troffaes and Dana Kelly)
2. Generalised Bayesian inference with sets of conjugate priors
   (joint work with Thomas Augustin)

Common-Cause Failures

Source: Wikimedia Commons, `http://commons.wikimedia.org/wiki/File:Fukushima_I_by_Digital_Globe.jpg`

Common-Cause Failures

Source: http://www.diakont.com/solutions/nuclear-energy/plant-systems/diesel-generator-control-systems/

Source: Wikimedia Commons, http://commons.wikimedia.org/wiki/File:Fukushima_I_by_Digital_Globe.jpg

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

common-cause failure

simultaneous failure of several redundant components due to a common or shared root cause (Høyland and Rausand 1994)

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

common-cause failure

simultaneous failure of several redundant components due to a common or shared root cause (Høyland and Rausand 1994)

- Reliability of redundant systems

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

common-cause failure

simultaneous failure of several redundant components due to a common or shared root cause (Høyland and Rausand 1994)

- Reliability of redundant systems
- Usually 2 – 4 emergency diesel generators per reactor

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

common-cause failure

simultaneous failure of several redundant components due to a common or shared root cause (Høyland and Rausand 1994)

- Reliability of redundant systems
- Usually 2 – 4 emergency diesel generators per reactor
- Sufficient cooling of core if one generator works

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

common-cause failure

simultaneous failure of several redundant components due to a common or shared root cause (Høyland and Rausand 1994)

- Reliability of redundant systems
- Usually 2 – 4 emergency diesel generators per reactor
- Sufficient cooling of core if one generator works
- Redundant components may not fail independently: common-cause failure

Common-Cause Failures

- All 12 generators (for 6 reactors) at Fukushima Daiichi were not available due to flooding of machine rooms (Tsunami caused by Tōhoku earthquake)

common-cause failure

simultaneous failure of several redundant components due to a common or shared root cause (Høyland and Rausand 1994)

- Reliability of redundant systems
- Usually 2 – 4 emergency diesel generators per reactor
- Sufficient cooling of core if one generator works
- Redundant components may not fail independently: common-cause failure

Must include common-cause failures in overall system reliability analysis

Common-Cause Failure Models

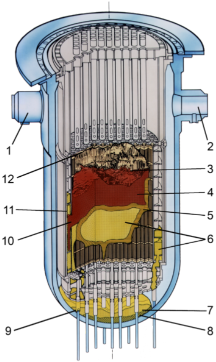

Above: CDC, `http://phil.cdc.gov/phil/` ID 1194

Right: Wikimedia Commons,
`http://commons.wikimedia.org/wiki/File:Graphic_TMI-2_Core_End-State_Configuration.png`

Basic Parameter Model

Basic Parameter Model (Mosleh et al. 1988)

- immediate repair
- failure events follow Poisson process
- system with k exchangeable components
- q_j: rate for failures involving *exact* j components ($j = 1, \ldots, k$)
- $(q_1, \ldots, q_k) =: \boldsymbol{q}$

$q_j \neq 0$ for $j \geq 2$: lack of independence for individual component failures

Basic Parameter Model

Basic Parameter Model (Mosleh et al. 1988)

- immediate repair
- failure events follow Poisson process
- system with k exchangeable components
- q_j: rate for failures involving *exact* j components ($j = 1, \ldots, k$)
- $(q_1, \ldots, q_k) =: \boldsymbol{q}$

$q_j \neq 0$ for $j \geq 2$: lack of independence for individual component failures

$\boldsymbol{q}$ is difficult to estimate directly:

- failure data often collected per component
- sparse data on joint failures

▶ reparametrisation: alpha-factor model

Alpha-Factor Model

Total Failure Rate

$$q_t = \sum_{j=1}^{k} \binom{k-1}{j-1} q_j \qquad (1)$$

total or *marginal failure rate:* failure rate obtained by looking just at single components

Alpha-Factors

$$\alpha_j = \frac{\binom{k}{j} q_j}{\sum_{\ell=1}^{k} \binom{k}{\ell} q_\ell} \qquad (2)$$

probability of j of the k components failing due to a common cause given that failure occurs

$$q_j = \frac{1}{\binom{k-1}{j-1}} \frac{j\alpha_j}{\sum_{\ell=1}^{k} \ell\alpha_\ell} q_t \qquad (3)$$

$$\boldsymbol{q} \Longleftrightarrow (q_t, \alpha_1, \ldots, \alpha_k)$$

Data

observed per-component failure rates to estimate q_t

Data

common-cause failure counts to estimate $(\alpha_1, \ldots, \alpha_k)$

Total Failure Rate: Data Model & Parameter Estimation

Poisson Process for Observed Per-Component Failures

$$p(M \mid q_t, T) = \frac{(q_t T)^M e^{-q_t T}}{M!} \tag{4}$$

where

- total failure rate q_t
- number of per-component (i.e. marginal) failures M := total number of component failures occured (two-component failure = two failures, ...)
- time under risk T := sum of time elapsed for each of the components

Estimation of q_t

usually immedially possible: use, e.g., maximum likelihood estimator

$$\hat{q}_t = \frac{M}{T} \tag{5}$$

Alpha-Factors: Data Model & Parameter Estimation

Multinomial Distribution for Common-Cause Failure Counts

$$p(\boldsymbol{n} \mid \boldsymbol{\alpha}) = \prod_{j=1}^{k} \alpha_j^{n_j} \tag{6}$$

where

- alpha-factor α_j := probability of j of the k components failing due to a common cause given that failure occurs
- failure count n_j := corresponding number of failures observed
- $\boldsymbol{n}$ denotes $(n_1, \ldots, n_k)$ and $\boldsymbol{\alpha}$ denotes $(\alpha_1, \ldots, \alpha_k)$

Estimation of $\boldsymbol{\alpha}$

maximum likelihood estimator:

$$\hat{\alpha}_j = \frac{n_j}{n}, \quad \text{where } \textstyle\sum_{j=1}^{n} n_j = n$$

Alpha-Factors: Data Model & Parameter Estimation

Estimation of α

maximum likelihood estimator:

$$\hat{\alpha}_j = \frac{n_j}{n}, \quad \text{where } \sum_{j=1}^{n} n_j = n$$

The Problem

- typically, for $j \geq 2$, the n_j are very low
  with zero being quite common for larger j
- zero counts = flat likelihoods ➡ $\hat{\alpha}_j = ?$

Alpha-Factors: Data Model & Parameter Estimation

Estimation of α

maximum likelihood estimator:

$$\hat{\alpha}_j = \frac{n_j}{n}, \quad \text{where } \sum_{j=1}^{n} n_j = n$$

The Problem

- typically, for $j \geq 2$, the n_j are very low
  with zero being quite common for larger j
- zero counts = flat likelihoods ➡ $\hat{\alpha}_j = ?$

➡ **need to rely on epistemic information: Bayesian inference**

Alpha-Factors: Data Model & Parameter Estimation

Estimation of α

maximum likelihood estimator:

$$\hat{\alpha}_j = \frac{n_j}{n}, \quad \text{where } \sum_{j=1}^{n} n_j = n$$

The Problem

- typically, for $j \geq 2$, the n_j are very low
  with zero being quite common for larger j
- zero counts = flat likelihoods ➡ $\hat{\alpha}_j = ?$

➡ **need to rely on epistemic information: Bayesian inference**

Bayesian inference procedure

prior + likelihood → posterior

using Bayes' Rule

All inferences are based on the posterior

Bayesian Inference: Dirichlet Prior

α considered as uncertain parameter on which we put...

Dirichlet Distribution (→ Dirichlet-Multinomial Model)

$$p(\alpha \mid s, t) \propto \prod_{j=1}^{k} \alpha_j^{s t - 1}$$

where (s, t) are *hyperparameters*

$$s > 0$$

$$t \in \Delta = \left\{ (t_1, \ldots, t_k) \colon t_1 \geq 0, \ldots, t_k \geq 0, \sum_{j=1}^{k} t_j = 1 \right\}$$

Bayesian Inference: Dirichlet Prior

α considered as uncertain parameter on which we put…

Dirichlet Distribution (→ Dirichlet-Multinomial Model)

$$p(\boldsymbol{\alpha} \mid s, \boldsymbol{t}) \propto \prod_{j=1}^{k} \alpha_j^{s\boldsymbol{t}-1}$$

where $(s, \boldsymbol{t})$ are *hyperparameters*

$$s > 0$$

$$\boldsymbol{t} \in \Delta = \left\{ (t_1, \ldots, t_k) : t_1 \geq 0, \ldots, t_k \geq 0, \sum_{j=1}^{k} t_j = 1 \right\}$$

Interpretation

- $\boldsymbol{t}$ = prior expectation of $\boldsymbol{\alpha}$, i.e., a prior guess for $\frac{n_j}{n}$, $j = 1, \ldots, n$
- s = determines spread and learning speed (see next slide)

Dirichlet Posterior

- posterior density for $\boldsymbol{\alpha}$ is again Dirichlet ($\rightarrow$ conjugacy):

$$p(\boldsymbol{\alpha} \mid \boldsymbol{n}, s, \boldsymbol{t}) \propto \prod_{j=1}^{k} \alpha_j^{st_j + n_j - 1} \tag{7}$$

- posterior expectation of α_j:

$$\mathrm{E}[\alpha_j \mid \boldsymbol{n}, s, \boldsymbol{t}] = \int_{\Delta} \alpha_j \, p(\boldsymbol{\alpha} \mid \boldsymbol{n}, s, \boldsymbol{t}) d\boldsymbol{\alpha} = \frac{s}{s+n} t_j + \frac{n}{s+n} \cdot \frac{n_j}{n} \tag{8}$$

we will focus on $\mathrm{E}[\alpha_j \mid \boldsymbol{n}, s, \boldsymbol{t}]$
(in a decision context, this expectation would typically end up in expressions for expected utility)

Example: Epistemic Information and Data

Example (from Kelly and Atwood 2011)

Consider a system with four redundant components ($k = 4$).
The analyst specifies the following prior expectation $\mu_{\text{spec},j}$ for each α_j:

$$\mu_{\text{spec},1} = 0.950 \quad \mu_{\text{spec},2} = 0.030 \quad \mu_{\text{spec},3} = 0.015 \quad \mu_{\text{spec},4} = 0.005 \tag{9}$$

We have 36 observations, in which 35 showed one component failing, and 1 showed two components failing:

$$n_1 = 35 \qquad n_2 = 1 \qquad n_3 = 0 \qquad n_4 = 0 \tag{10}$$

Non-Informative Priors

large variation in posterior under different non-informative priors

- with constrained maximum entropy prior (Atwood 1996; Kelly and Atwood 2011):

$$\mathrm{E}[\alpha_1 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.967 \qquad \mathrm{E}[\alpha_2 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.028$$
$$\mathrm{E}[\alpha_3 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.003 \qquad \mathrm{E}[\alpha_4 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.001$$

- with uniform prior $t_j = 0.25$ and $s = 4$:

$$\mathrm{E}[\alpha_1 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.9 \qquad \mathrm{E}[\alpha_2 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.05$$
$$\mathrm{E}[\alpha_3 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.025 \qquad \mathrm{E}[\alpha_4 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.025$$

- with Jeffreys' prior $t_j = 0.25$ and $s = 2$:

$$\mathrm{E}[\alpha_1 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.9342 \qquad \mathrm{E}[\alpha_2 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.0395$$
$$\mathrm{E}[\alpha_3 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.0132 \qquad \mathrm{E}[\alpha_4 \mid \boldsymbol{n}, s, \boldsymbol{t}] = 0.0132$$

Imprecise Dirichlet Model: Definition

Troffaes, Walter, and Kelly (2014):
model vague prior info more cautiously

Imprecise Dirichlet Model (IDM) for Common-Cause Failure

use a set of hyperparameters (Walley 1991; Walley 1996)

$$\mathcal{H} = \Big\{(s, \boldsymbol{t})\colon s \in [\underline{s}, \overline{s}],\ \boldsymbol{t} \in \Delta,\ t_j \in [\underline{t}_j, \overline{t}_j]\Big\}$$

Imprecise Dirichlet Model: Definition

Troffaes, Walter, and Kelly (2014):
model vague prior info more cautiously

Imprecise Dirichlet Model (IDM) for Common-Cause Failure

use a set of hyperparameters (Walley 1991; Walley 1996)

$$\mathcal{H} = \left\{(s, \boldsymbol{t}) \colon s \in [\underline{s}, \overline{s}],\ \boldsymbol{t} \in \Delta,\ t_j \in [\underline{t}_j, \overline{t}_j]\right\}$$

Interpretation

- we are doing a sensitivity analysis (á la robust Bayes) over $(s, \boldsymbol{t}) \in \mathcal{H}$
- we take a set of priors based on $\mathcal{H}$ as model for prior information (details later)

Imprecise Dirichlet Model: Definition

Troffaes, Walter, and Kelly (2014):
model vague prior info more cautiously

Imprecise Dirichlet Model (IDM) for Common-Cause Failure

use a set of hyperparameters (Walley 1991; Walley 1996)

$$\mathcal{H} = \left\{(s, \boldsymbol{t}) \colon s \in [\underline{s}, \overline{s}],\ \boldsymbol{t} \in \Delta,\ t_j \in [\underline{t}_j, \overline{t}_j]\right\}$$

Interpretation

- we are doing a sensitivity analysis (á la robust Bayes) over $(s, \boldsymbol{t}) \in \mathcal{H}$
- we take a set of priors based on $\mathcal{H}$ as model for prior information (details later)

Analyst has to specify ('elicit')
bounds $[\underline{s}, \overline{s}]$ and bounds $[\underline{t}_j, \overline{t}_j]$ for each $j \in \{1, \ldots, k\}$

Imprecise Dirichlet Model: Elicitation

- $[\underline{t}_j, \overline{t}_j]$? Cautious interpretation of prior specifications $\mu_{\text{spec},j}$:

$$[\underline{t}_1, \overline{t}_1] = [0.950, 1] \qquad [\underline{t}_2, \overline{t}_2] = [0, 0.030]$$

$$[\underline{t}_3, \overline{t}_3] = [0, 0.015] \qquad [\underline{t}_4, \overline{t}_4] = [0, 0.005]$$

Imprecise Dirichlet Model: Elicitation

- $[\underline{t}_j, \overline{t}_j]$? Cautious interpretation of prior specifications $\mu_{\text{spec},j}$:

$$[\underline{t}_1, \overline{t}_1] = [0.950, 1] \qquad [\underline{t}_2, \overline{t}_2] = [0, 0.030]$$
$$[\underline{t}_3, \overline{t}_3] = [0, 0.015] \qquad [\underline{t}_4, \overline{t}_4] = [0, 0.005]$$

- $[\underline{s}, \overline{s}]$? Good (1965):

  reason about posterior expectations for hypothetical data

Imprecise Dirichlet Model: Elicitation

- $[\underline{t}_j, \overline{t}_j]$? Cautious interpretation of prior specifications $\mu_{\text{spec},j}$:

$$[\underline{t}_1, \overline{t}_1] = [0.950, 1] \qquad [\underline{t}_2, \overline{t}_2] = [0, 0.030]$$
$$[\underline{t}_3, \overline{t}_3] = [0, 0.015] \qquad [\underline{t}_4, \overline{t}_4] = [0, 0.005]$$

- $[\underline{s}, \overline{s}]$? Good (1965):

  reason about posterior expectations for hypothetical data

$\overline{s}$ = number of one-component failures required
to reduce the upper probabilities of multi-component failure by half

$\underline{s}$ = number of multi-component failures required
to reduce the lower probability of one-component failure by half

Imprecise Dirichlet Model: Elicitation

$\overline{s}$ = number of one-component failures required
to reduce the upper probabilities of multi-component failure by half

$\underline{s}$ = number of multi-component failures required
to reduce the lower probability of one-component failure by half

Imprecise Dirichlet Model: Elicitation

$\overline{s}$ = number of one-component failures required
to reduce the upper probabilities of multi-component failure by half

$\underline{s}$ = number of multi-component failures required
to reduce the lower probability of one-component failure by half

Reasonable values in example:

- $\overline{s} = 10$: after observing 10 one-component failures
  ➡ halve upper probabilities of multi-component failures
- $\underline{s} = 1$: immediate multi-component failure
  ➡ keen to reduce lower probability for one-component failure

Imprecise Dirichlet Model: Elicitation

$\overline{s}$ = number of one-component failures required
to reduce the upper probabilities of multi-component failure by half

$\underline{s}$ = number of multi-component failures required
to reduce the lower probability of one-component failure by half

Reasonable values in example:

- $\overline{s} = 10$: after observing 10 one-component failures
  ➡ halve upper probabilities of multi-component failures
- $\underline{s} = 1$: immediate multi-component failure
  ➡ keen to reduce lower probability for one-component failure

Difference between $\underline{s}$ and $\overline{s}$ reflects a level of caution:
The rate at which we reduce upper probabilities
is less than the rate at which we reduce lower probabilities

Imprecise Dirichlet Model: Inference

With $[\underline{s}, \overline{s}] = [1, 10]$, we get...

prior bounds + data → posterior bounds

j	$\underline{t}_j$	$\overline{t}_j$	n_j	$\underline{E}[\alpha_j \mid \boldsymbol{n}, \mathcal{H}]$	$\overline{E}[\alpha_j \mid \boldsymbol{n}, \mathcal{H}]$
1	0.950	1	35	0.967	0.978
2	0	0.030	1	0.0270	0.0283
3	0	0.015	0	0	0.00326
4	0	0.005	0	0	0.00109

Gamma Prior and Posterior

q_t considered as uncertain parameter on which we put...

Gamma Distribution

$$p(q_t \mid u, v) \propto q_t^{uv-1} e^{-q_t u} \tag{11}$$

where (u, v) are hyperparameters with $u > 0$ and $v > 0$.

Interpretation

- v = prior expectation of q_t
- u = determines learning speed (just like s in the IDM)

- posterior density for q_t is again Gamma:

$$p(q_t \mid M, T, u, v) \propto q_t^{uv+M-1} e^{-q_t(u+T)} \tag{12}$$

- posterior expectation of q_t:

$$\mathrm{E}[q_t \mid M, T, u, v] = \frac{u}{T+u} v + \frac{T}{T+u} \cdot \frac{M}{T} \tag{13}$$

Imprecise Gamma Model

use a set of hyperparameters:

$$= \Big\{ (u, v) : u \in [\underline{u}, \overline{u}],\ v \in [\underline{v}, \overline{v}] \Big\} \qquad (14)$$

- $[\underline{v}, \overline{v}]$? Bounds for prior expectation of q_t should be easy to find (choosing $\underline{v} = 0$ is possible)
- $[\underline{u}, \overline{u}]$? Similar reasoning as for the IDM leads to...

$\overline{u}$ = timespan for observing the process required to raise the lower expectation of q_t from 0 to half of observed failure rate $\frac{M}{T}$ ($\underline{v} = 0$ is assumed)

$\underline{u}$ = timespan for observing the process *without any failures* required to reduce the lower expectation of q_t by half ($\underline{v} > 0$ is assumed)

$\underline{u} = \overline{u}$ can be reasonable here, as zero counts are less of an issue

Inference on Common-Cause Failure Rates q_j

combine our models for α and q_t by using Eq. (3):

$$q_j = g_j(\alpha) q_t \qquad \text{where} \qquad g_j(\alpha) = \frac{1}{\binom{k-1}{j-1}} \frac{j\alpha_j}{\sum_{\ell=1}^{k} \ell \alpha_\ell}$$

The Problem

no closed expression for $\mathrm{E}[g_j(\alpha) \mid \dots]$ due to rational function of α

The Good News

naive approximation $\tilde{g}_j(\alpha)$ of $g_j(\alpha)$ by Taylor expansion works surprisingly well (absolute error term available)

$$\mathrm{E}[q_j \mid \boldsymbol{n}, s, \boldsymbol{t}; M, T, u, v] \approx \mathrm{E}\left[\tilde{g}_j(\alpha) \mid \boldsymbol{n}, s, \boldsymbol{t}\right] \mathrm{E}[q_t \mid M, T, u, v] \qquad (15)$$

(q_t and α are assumed to be independent)

Global Sensitivity Analysis

We can do a **global sensitivity analysis** for $\mathrm{E}[q_j \mid \ldots]$

➡ bounds for $\mathrm{E}[q_j \mid \ldots]$ taking into account approximation error and *epistemic uncertainty expressed through*   *and*   :

$$\underline{\mathrm{E}}[q_j \mid \boldsymbol{n}, M, T, \quad , \quad] \approx \underline{\mathrm{E}}[\tilde{g}_j(\alpha) \mid \boldsymbol{n}, \quad]\, \underline{\mathrm{E}}[q_t \mid M, T, \quad] \tag{16}$$

where

$$\underline{\mathrm{E}}[\tilde{g}_j(\alpha) \mid \boldsymbol{n}, \quad] = \min_{(s,\boldsymbol{t}) \in} \mathrm{\underline{E}}[\tilde{g}_j(\alpha) \mid \boldsymbol{n}, s, \boldsymbol{t}] \quad \text{(by num. optimization)} \tag{17}$$

$$\underline{\mathrm{E}}[q_t \mid M, T, \quad] = \min_{(u,v) \in} \mathrm{E}[q_t \mid M, T, u, v] \quad \text{(by closed form solution)} \tag{18}$$

Do the same for $\overline{\mathrm{E}}[q_j \mid \boldsymbol{n}, M, T, \quad , \quad]$ by replacing min with max.

Intermediate Summary

- bounds, rather than precise values, are desirable due to inferences being strongly sensitive to the prior particularly when faced with zero counts.

Intermediate Summary

- bounds, rather than precise values, are desirable due to inferences being strongly sensitive to the prior particularly when faced with zero counts.
- simple ways to elicit the parameters of the model by reasoning on hypothetical data

Intermediate Summary

- bounds, rather than precise values, are desirable due to inferences being strongly sensitive to the prior particularly when faced with zero counts.
- simple ways to elicit the parameters of the model by reasoning on hypothetical data
- sets of hyperparameters allow a full sensitivity analysis reflecting epistemic uncertainty on all parts of the model

Intermediate Summary

- bounds, rather than precise values, are desirable due to inferences being strongly sensitive to the prior particularly when faced with zero counts.
- simple ways to elicit the parameters of the model by reasoning on hypothetical data
- sets of hyperparameters allow a full sensitivity analysis reflecting epistemic uncertainty on all parts of the model
- use variance instead of bounds on expectations? (operational interpretation of variance of unknown parameter versus direct bounds on expectation of unknown parameter)

Intermediate Summary

- bounds, rather than precise values, are desirable
  due to inferences being strongly sensitive to the prior
  particularly when faced with zero counts.
- simple ways to elicit the parameters of the model
  by reasoning on hypothetical data
- sets of hyperparameters allow a full sensitivity analysis
  reflecting epistemic uncertainty on all parts of the model
- use variance instead of bounds on expectations?
  (operational interpretation of variance of unknown parameter
  versus direct bounds on expectation of unknown parameter)
- hyperparameter sets have very specific form
  (chosen to allow easy elicitation),
  how does shape influence posterior inferences?

Intermediate Summary

- bounds, rather than precise values, are desirable due to inferences being strongly sensitive to the prior particularly when faced with zero counts.
- simple ways to elicit the parameters of the model by reasoning on hypothetical data
- sets of hyperparameters allow a full sensitivity analysis reflecting epistemic uncertainty on all parts of the model
- use variance instead of bounds on expectations? (operational interpretation of variance of unknown parameter versus direct bounds on expectation of unknown parameter)
- hyperparameter sets have very specific form (chosen to allow easy elicitation), how does shape influence posterior inferences?
- is it possible to generalise this method to other problems?

Canonical Conjugate Priors

Multinomial, Poisson are examples for a canonical exponential family:

$(x_1, \dots, x_n) = \boldsymbol{x} \overset{iid}{\sim}$ canonical exponential family

$$p(\boldsymbol{x} \mid \theta) \propto \exp\left\{\langle\psi, \tau(\boldsymbol{x})\rangle - n\mathrm{b}(\psi)\right\} \qquad \left[\psi \text{ transformation of } \theta\right] \qquad (19)$$

(includes also Binomial, Normal, Exponential, Dirichlet, Gamma, ...)

- conjugate prior: $p(\psi \mid n^{(0)}, y^{(0)}) \propto \exp\left\{n^{(0)}\left[\langle\psi, y^{(0)}\rangle - \mathrm{b}(\psi)\right]\right\}$
- (conjugate) posterior: $p(\psi \mid n^{(0)}, y^{(0)}, \boldsymbol{x}) \propto \exp\left\{n^{(n)}\left[\langle\psi, y^{(n)}\rangle - \mathrm{b}(\psi)\right]\right\}$

where $y^{(n)} = \dfrac{n^{(0)}}{n^{(0)} + n} \cdot y^{(0)} + \dfrac{n}{n^{(0)} + n} \cdot \dfrac{\tau(\boldsymbol{x})}{n}$ and $n^{(n)} = n^{(0)} + n$

Interpretation

- $n^{(0)}$ = determines spread and learning speed
- $y^{(0)}$ = prior expectation of $\tau(\boldsymbol{x})/n$

Imprecision

Add imprecision as new modelling dimension:
Sets of priors model uncertainty in probability statements and allow to better model partial or vague information on θ

Imprecision

Add imprecision as new modelling dimension:
Sets of priors model uncertainty in probability statements and allow to better model partial or vague information on θ

Interpretation

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
→ $P(\text{win}) = [1/100, 7/100]$

Bayesian Inference with Sets of Conjugate Priors

Standard Bayesian inference procedure

$$\text{prior} + \text{likelihood} \to \text{posterior}$$

using Bayes' Rule

All inferences are based on the posterior

(e.g., point estimate $\mathrm{E}[\psi \mid \boldsymbol{x}, n^{(0)}, y^{(0)}] = \mathrm{E}[\psi \mid n^{(n)}, y^{(n)}]$)

Bayesian Inference with Sets of Conjugate Priors

Standard Bayesian inference procedure

prior + likelihood → posterior

using Bayes' Rule

All inferences are based on the posterior

(e.g., point estimate $\mathrm{E}[\psi \mid \boldsymbol{x}, n^{(0)}, y^{(0)}] = \mathrm{E}[\psi \mid n^{(n)}, y^{(n)}]$)

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

Bayesian Inference with Sets of Conjugate Priors

Standard Bayesian inference procedure

prior + likelihood → posterior

using Bayes' Rule

All inferences are based on the posterior

(e.g., point estimate $\mathrm{E}[\psi \mid \boldsymbol{x}, n^{(0)}, y^{(0)}] = \mathrm{E}[\psi \mid n^{(n)}, y^{(n)}]$)

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

Generalised Bayesian inference procedure

set of priors + likelihood → set of posteriors

All inferences are based on the set of posteriors

(e.g., $\underline{\mathrm{E}}[\psi \mid \boldsymbol{x}, \Pi^{(0)}], \overline{\mathrm{E}}[\psi \mid \boldsymbol{x}, \Pi^{(0)}]$)

Coherence (consistency of inferences) ensured by using

Generalised Bayes' Rule (GBR, Walley 1991)

= element-wise application of Bayes' Rule

General Model Properties

Model framework has favourable inference properties
(see Walter 2013, §3.1) and is very easy to handle:

- Hyperparameter set defines set of priors $\mathcal{M}^{(0)}$

General Model Properties

Model framework has favourable inference properties (see Walter 2013, §3.1) and is very easy to handle:

- Hyperparameter set defines set of priors $\mathcal{M}^{(0)}$
- Due to conjugacy, set of posteriors $\mathcal{M}^{(n)}$ defined by updated hyperparameter set

General Model Properties

Model framework has favourable inference properties (see Walter 2013, §3.1) and is very easy to handle:

- Hyperparameter set defines set of priors $\mathcal{M}^{(0)}$
- Due to conjugacy, set of posteriors $\mathcal{M}^{(n)}$ defined by updated hyperparameter set
- $\rightarrow$ is easy:

$$n^{(n)} = n^{(0)} + n \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n} y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{\tau(\boldsymbol{x})}{n} \tag{20}$$

General Model Properties

Model framework has favourable inference properties (see Walter 2013, §3.1) and is very easy to handle:

- Hyperparameter set defines set of priors $\mathcal{M}^{(0)}$
- Due to conjugacy, set of posteriors $\mathcal{M}^{(n)}$ defined by updated hyperparameter set
- $\rightarrow$ is easy:

$$n^{(n)} = n^{(0)} + n \qquad y^{(n)} = \frac{n^{(0)}}{n^{(0)} + n} y^{(0)} + \frac{n}{n^{(0)} + n} \cdot \frac{\tau(\boldsymbol{x})}{n} \tag{20}$$

- Often, optimising over $(n^{(n)}, \boldsymbol{y}^{(n)}) \in$ is also easy: closed form solution for $y^{(n)}$ = posterior 'guess' for $\frac{\tau(\boldsymbol{x})}{n}$ given has 'nice' shape (as used in the common-cause model)

Parameter Set Shapes

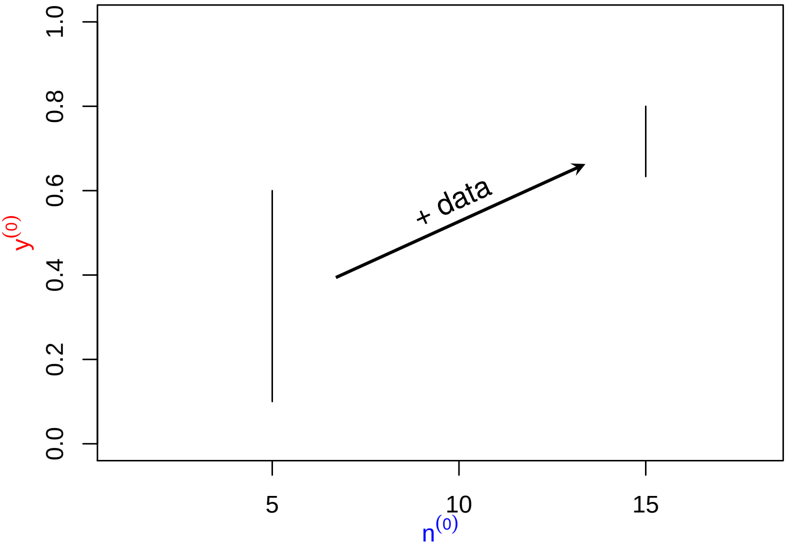

Parameter Set Shapes

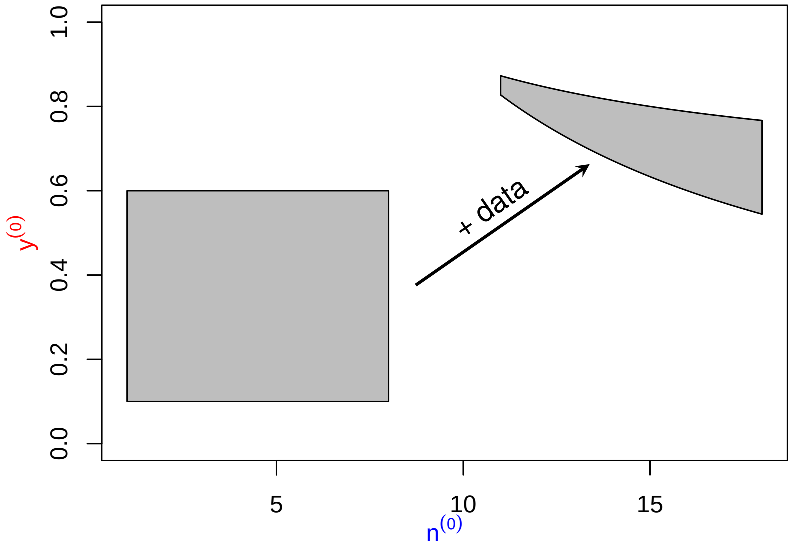

Parameter Set Shapes

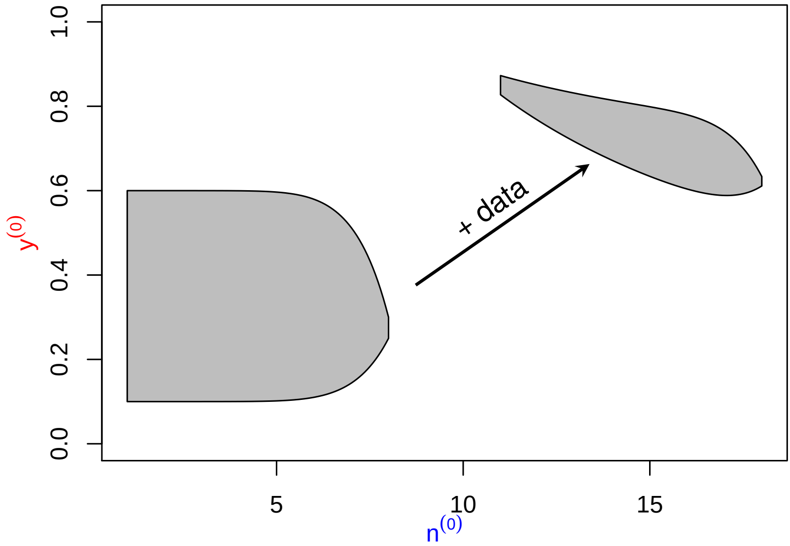

Parameter Set Shapes & Prior-Data Conflict

Prior-Data Conflict

- *informative prior beliefs* and *trusted data* (sampling model correct, no outliers, etc.) are in conflict
- "[. . .] the prior [places] its mass primarily on distributions in the sampling model for which the observed data is surprising" (Evans and Moshonov 2006)
- there are not enough data to overrule the prior

Parameter Set Shapes & Prior-Data Conflict

Prior-Data Conflict

- *informative prior beliefs* and *trusted data* (sampling model correct, no outliers, etc.) are in conflict
- "[. . .] the prior [places] its mass primarily on distributions in the sampling model for which the observed data is surprising" (Evans and Moshonov 2006)
- there are not enough data to overrule the prior

The Problem

Many Bayesian models are insensitive to prior-data conflict!

Scaled Normal Data $\boldsymbol{x} \overset{iid}{\sim} \mathrm{N}(\mu, 1)$: $\mu \sim \mathrm{N}(y^{(0)}, 1/n^{(0)})$

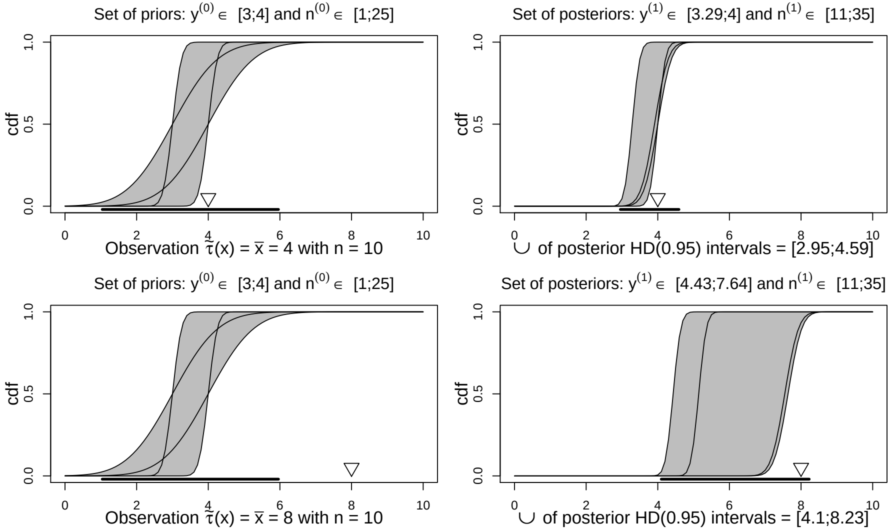

Conclusion

- Conjugate priors are a convenient tool for Bayesian inference but there are some pitfalls
  - Hyperparameters $n^{(0)}, y^{(0)}$ are easy to interpret and elicit
  - Averaging property makes calculations simple, but leads to inadequate model behaviour in case of prior-data conflict

Conclusion

- Conjugate priors are a convenient tool for Bayesian inference but there are some pitfalls
  - Hyperparameters $n^{(0)}, y^{(0)}$ are easy to interpret and elicit
  - Averaging property makes calculations simple, but leads to inadequate model behaviour in case of prior-data conflict
- Sets of conjugate priors maintain advantages & mitigate issues
  - Hyperparameter set shape is important
  - Reasonable choice: *rectangular* $= [\underline{n}^{(0)}, \overline{n}^{(0)}] \times [\underline{y}^{(0)}, \overline{y}^{(0)}]$ (Walter & Augustin 2009: *generalised iLUCK-models*, `luck`)
  - Bounds for hyperparameters are easy to interpret and elicit
  - Additional imprecison in case of prior-data conflict leads to cautious inferences if, and only if, caution is needed
  - Shape for more precision in case of strong prior-data agreement is in development (joint work with Frank Coolen and Mik̡ Bickis)

References

Atwood, C. (1996). “Constrained noninformative priors in risk assessment”. In: *Reliability Engineering and System Safety* 53, pp. 37–46. DOI: 10.1016/0951-8320(96)00026-9.

Evans, M. and H. Moshonov (2006). “Checking for Prior-Data Conflict”. In: *Bayesian Analysis* 1, pp. 893–914. URL: http://projecteuclid.org/euclid.ba/1340370946.

Good, I. (1965). *The estimation of probabilities*. Cambridge (MA): MIT Press.

Høyland, Arnljot and Marvin Rausand (1994). *System reliability theory: models and statistical methods*. New York (NY): Wiley.

Kelly, Dana and Corwin Atwood (2011). “Finding a minimally informative Dirichlet prior distribution using least squares”. In: *Reliability Engineering and System Safety* 96.3, pp. 398–402. DOI: 10.1016/j.ress.2010.11.008.

Mosleh, A. et al. (1988). *Procedures for treating common cause failures in safety and reliability studies: Procedural framework and examples*. Tech. rep. NUREG/CR-4780. PLG Inc.

Quaeghebeur, E. and G. de Cooman (2005). “Imprecise probability models for inference in exponential families”. In: *ISIPTA ’05*. Ed. by F. Cozman, R. Nau, and T. Seidenfeld. Manno: SIPTA, pp. 287–296.

Troffaes, Matthias C. M., Gero Walter, and Dana L. Kelly (2014). “A Robust Bayesian Approach to Modelling Epistemic Uncertainty in Common-Cause Failure Models”. In: *Reliability Engineering & System Safety* 125, pp. 13–21. DOI: 10.1016/j.ress.2013.05.022.

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

Walley, P. (1996). “Inferences from multinomial data: Learning about a bag of marbles”. In: *Journal of the Royal Statistical Society, Series B* 58.1, pp. 3–34.

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.