Actuaries Institute.

Life Insurance Applications of Bayesian Models

Colin Yellowlees and Tikiri Appuhamy

May 2026

Presented at the 2026 All Actuaries Summit

What are Bayesian models?

- Bayesian modelling in simple terms involve updating your existing assumptions based on new information. The existing assumptions are referred to as prior distribution and the assumptions after Bayesian updating are known as the posterior distribution.

- They update these assumptions through Bayes theorem:

$$\overbrace{\Pr(\,\theta \mid Y\,)}^{\text{Posterior}} = \frac{\overbrace{\Pr(\theta)}^{\text{Prior}} \times \overbrace{\Pr(Y \mid \theta)}^{\text{Likelihood}}}{\Pr(Y)}$$

- The key difference in Bayesian modelling is that practitioners need to assign a distribution for each prior assumption. Judgement is required to set priors.

Why use Bayesian models?

Captures the uncertainty in assumptions and the statistical process

- Actuarial models contain a number of assumptions – traditional methods fail to capture assumption uncertainty.
- Bayesian models are a systematic way to measure assumption uncertainty.

Naturally generative

- Both Actuarial and Bayesian models attempt to model the data generating process hence Bayesian model fit naturally into the Actuarial toolbox.

Suitable for small-medium datasets

- When priors are selected carefully, Bayesian models can regularise assumption to prevent overfitting.
- Priors can also be used allow for expert judgement or existing assumptions.

Overconfidence and point estimates

- Traditional Actuarial practice focuses on the development of point estimates assumptions e.g. decline rates, incidence rates, claims development patterns.

- Each point estimate has uncertainty that is lost when they are ultimately used in pricing and valuation processes.

- We attempt recover a measurement of assumption uncertainty when we think about how to measure capital or valuing experience-sharing mechanisms.

- A very simple example of this can be found in group pricing where we follow the following process:

Decline Rate → IBNR → Loss ratio → Pricing Margins

Overconfidence and point estimates – IBNR & decline rates

There is a range of possible assumptions that could be chosen when we conduct our analysis, however traditional analysis forces us to choose only one set of assumptions!

Event Year	0:1	1:2	2:3	3:4	4:5	5:6	6:7	7:8	8:9
2016	3.79	1.58	1.35	1.18	1.15	1.08	1.07	1.04	1.02
2017	2.87	2.18	1.24	1.22	1.13	1.09	1.03	1.02	
2018	2.23	1.72	1.31	1.18	1.14	1.08	1.04		
2019	4.89	1.62	1.25	1.18	1.16	1.02			
2020	2.90	1.89	1.34	1.21	1.16				
2021	2.29	1.97	1.27	1.12					
2022	3.52	2.31	1.50						
2023	2.57	1.37							
2024	2.78								
2025									
Selected Assumption	2.91	1.78	1.32	1.18	1.15	1.07	1.05	1.03	1.02

Large variability

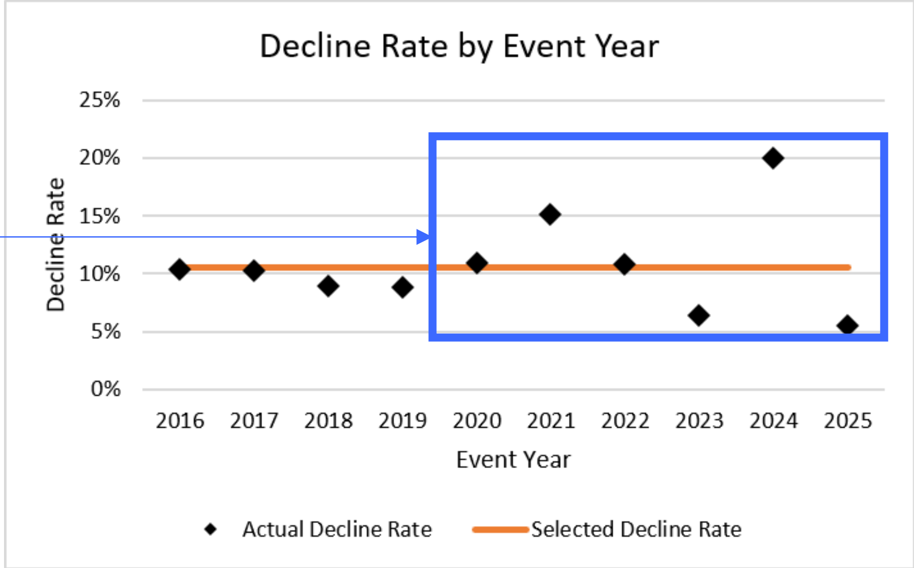

Figure 1: Age to age claims development factors

Figure 2: Decline rate analysis

Overconfidence and point estimates – loss ratios

- Based on the decline rate and claims development pattern selected in the previously slides lead to the following loss ratios.
- We've lost information on the uncertainty inherent in the IBNR and RBNA!

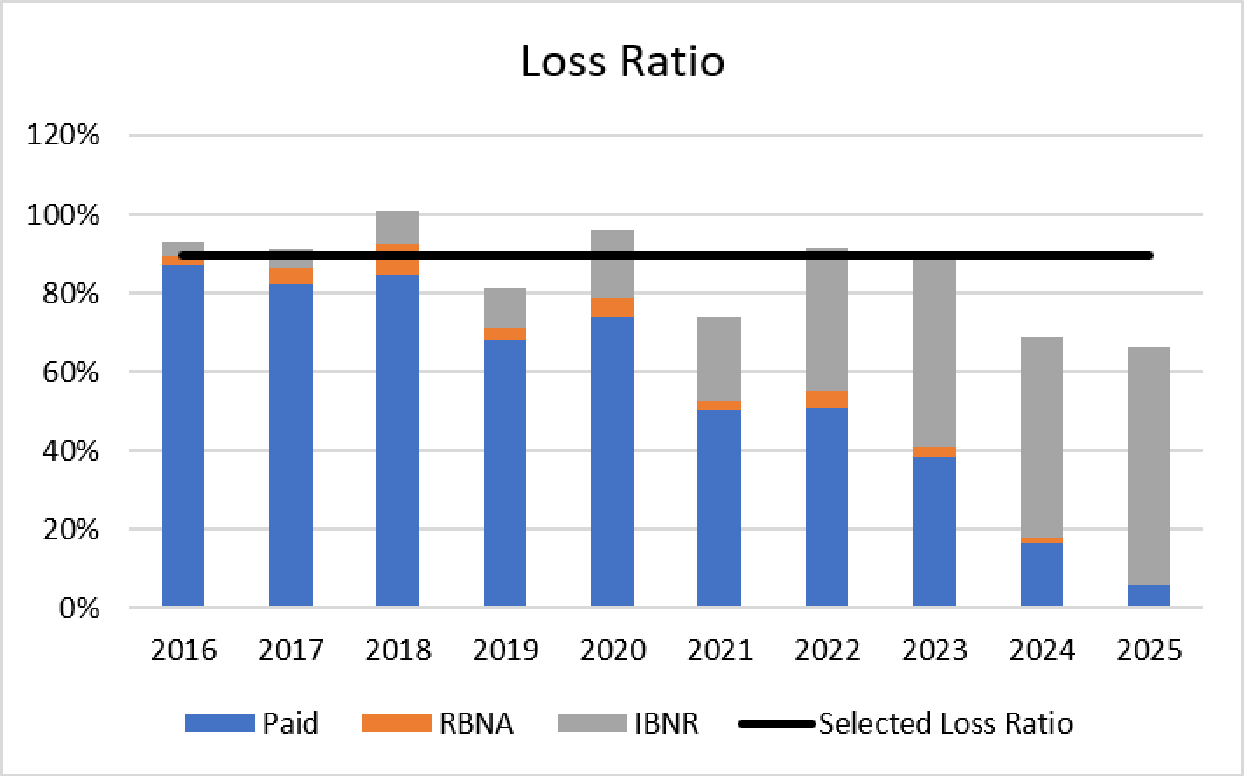

Figure 1: Loss ratios by event year

Overconfidence and point estimates – Bayesian assumptions

Bayesian model captures
the parameter uncertainty

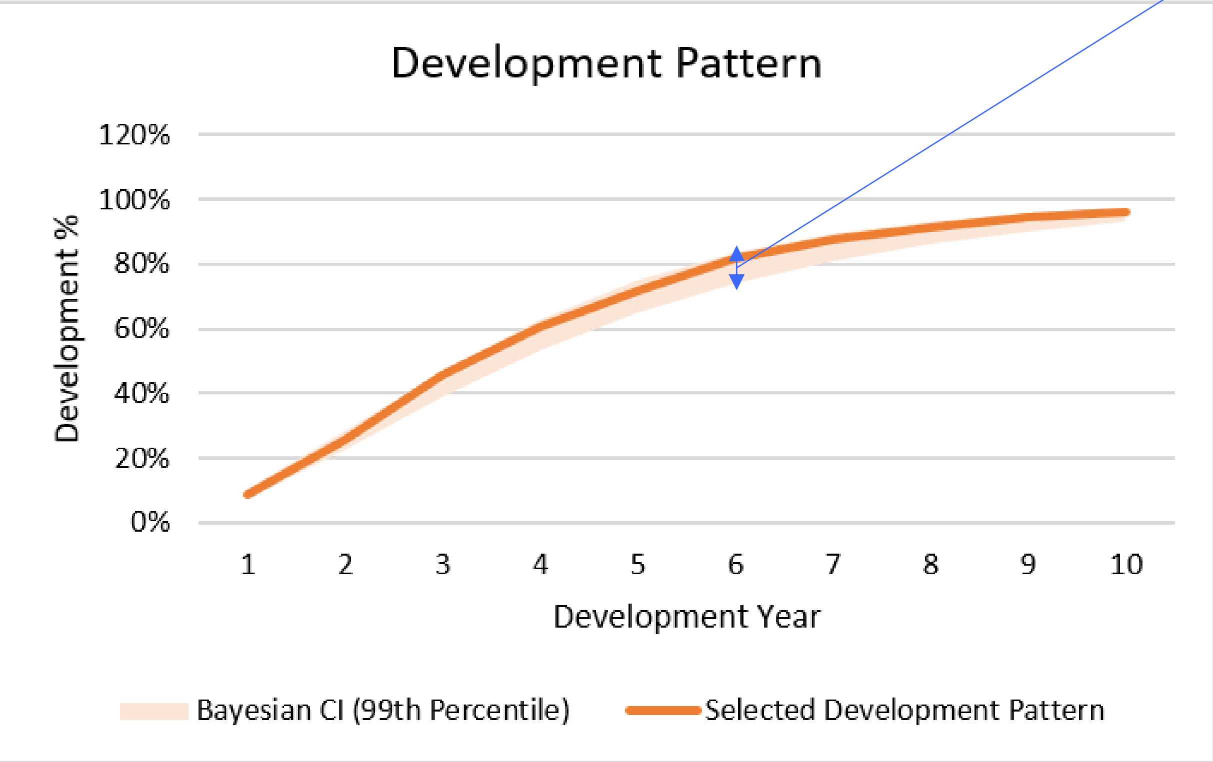

Figure 1: Bayesian IBNR cumulative development factors

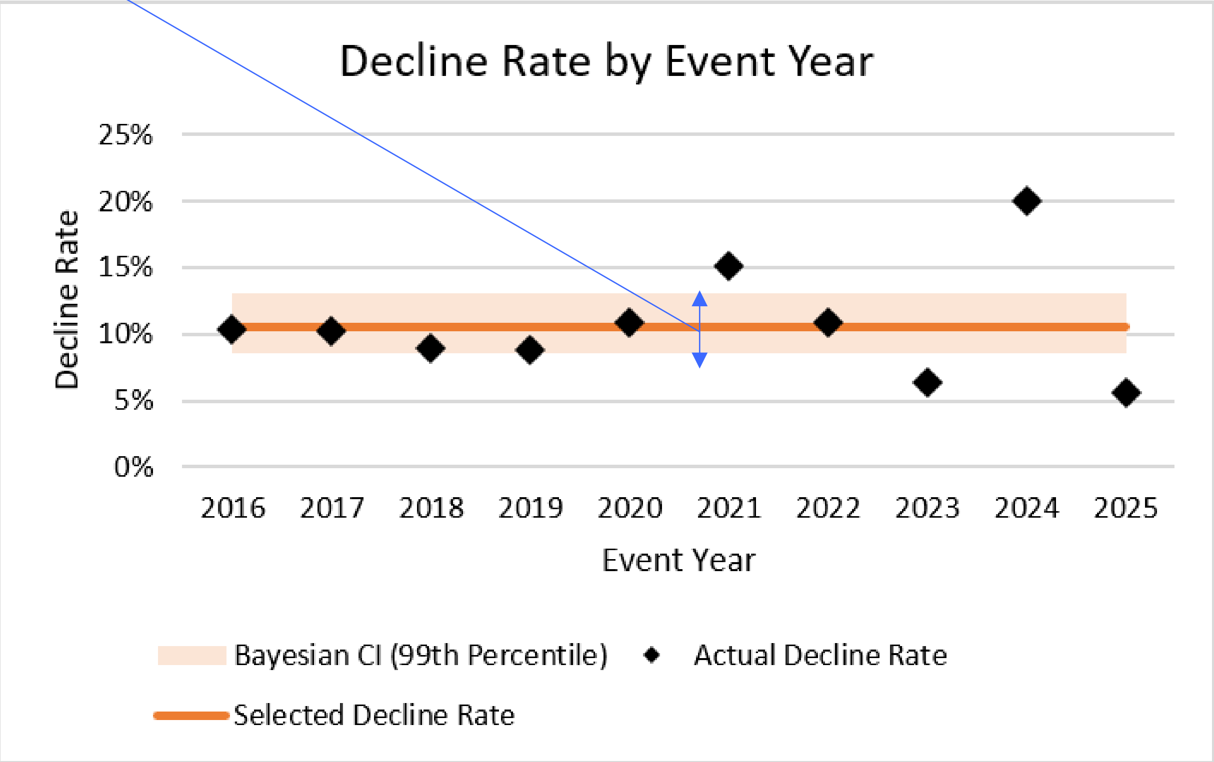

Figure 2: Bayesian decline rate

Overconfidence and point estimates – Bayesian assumptions

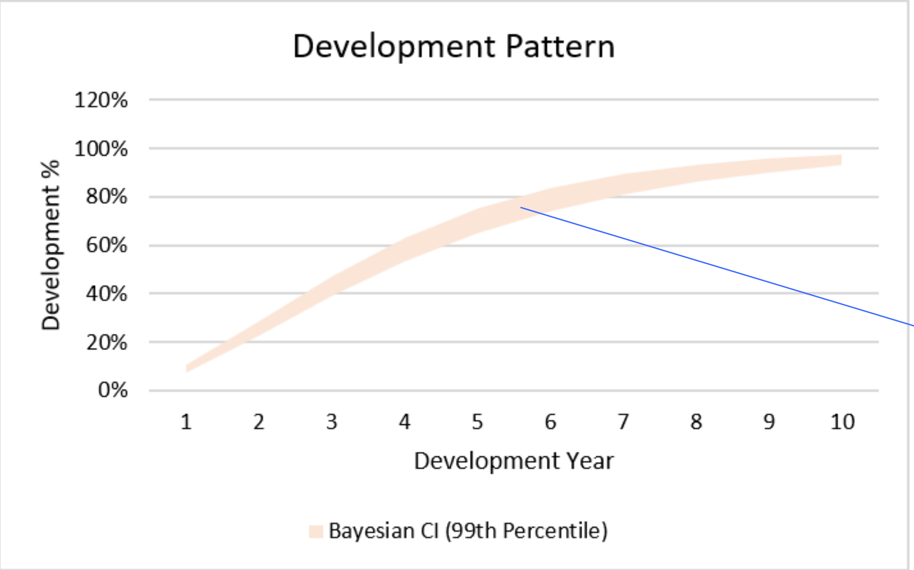

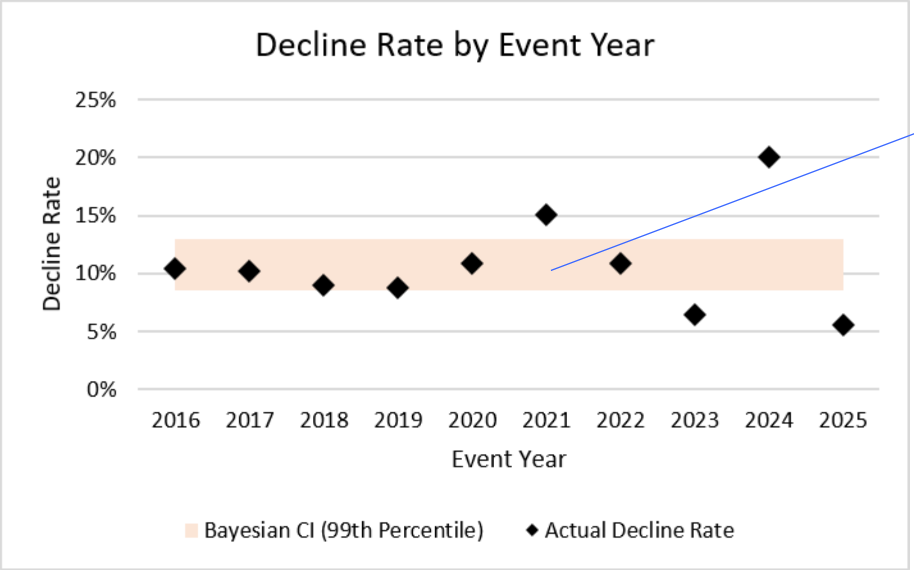

The Bayesian model captures the uncertainty in the development pattern and the decline rate assumptions in the loss ratio.

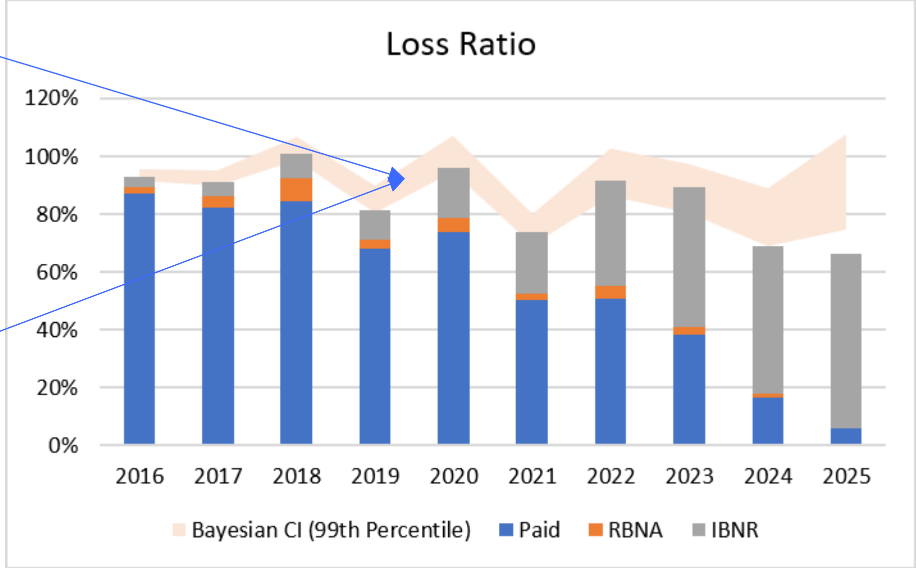

Overconfidence and point estimates – misestimation risk and process risk

To measure the process risk, we assume that the claims follow a compound Poisson process where the size of claims are log-normally distributed.

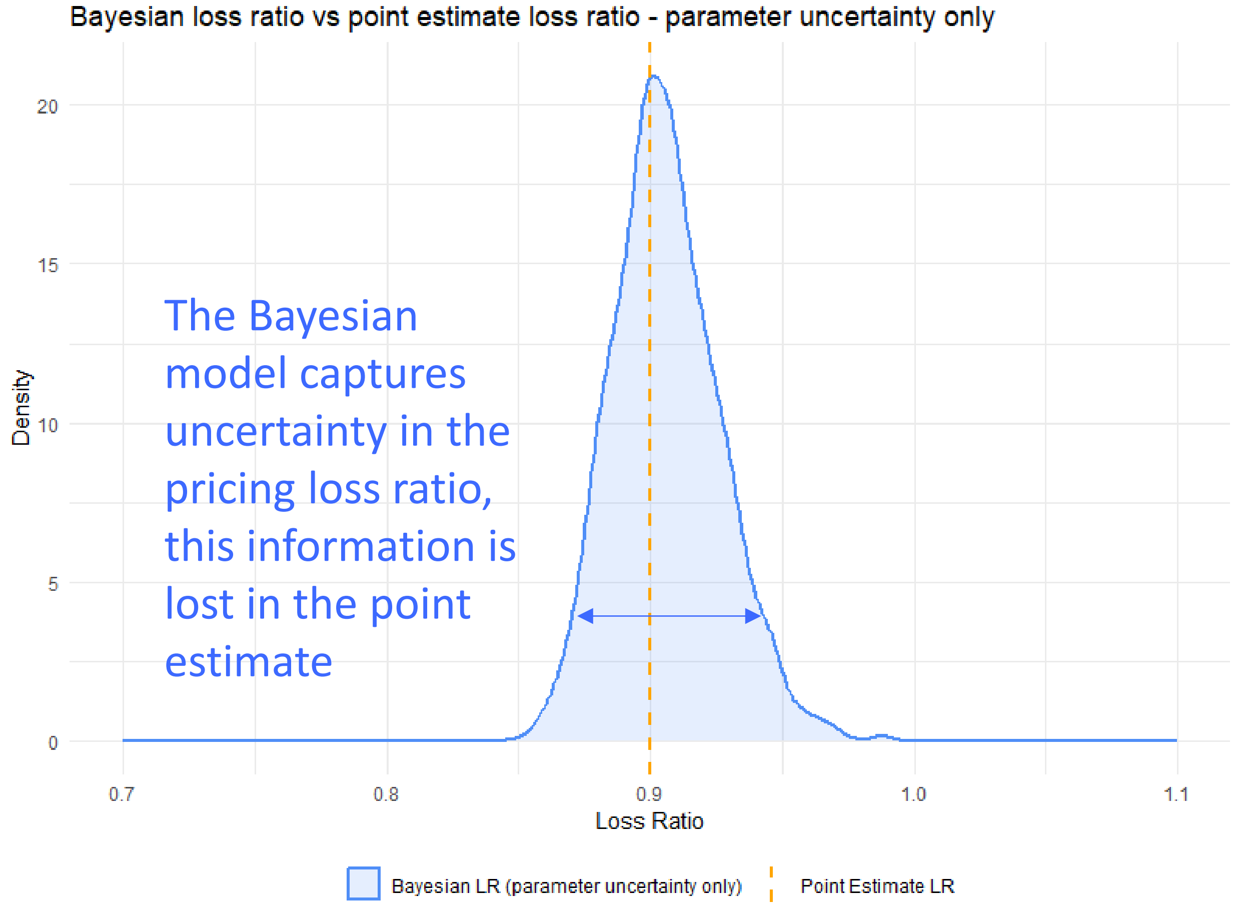

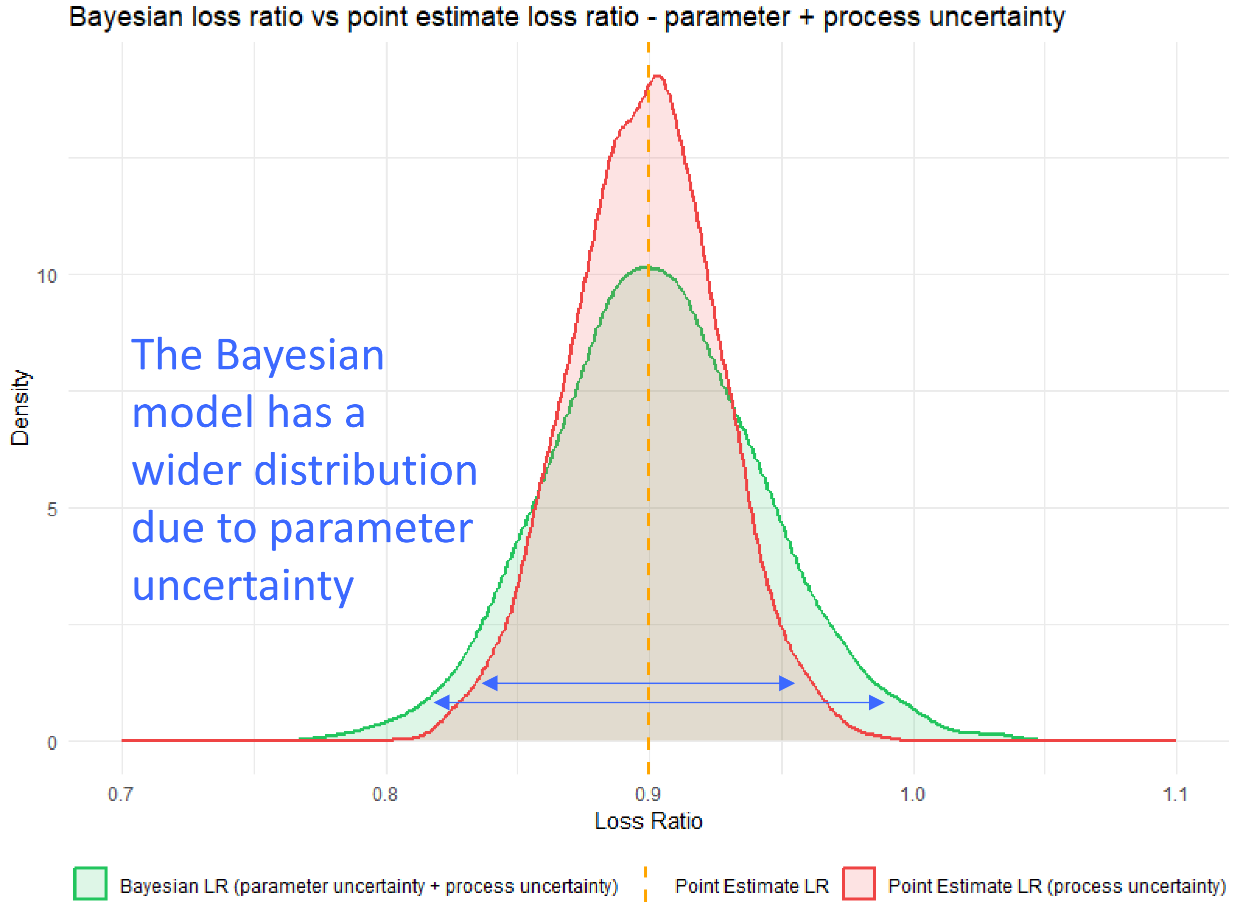

Bayesian workflow

Specific to Bayesian modelling

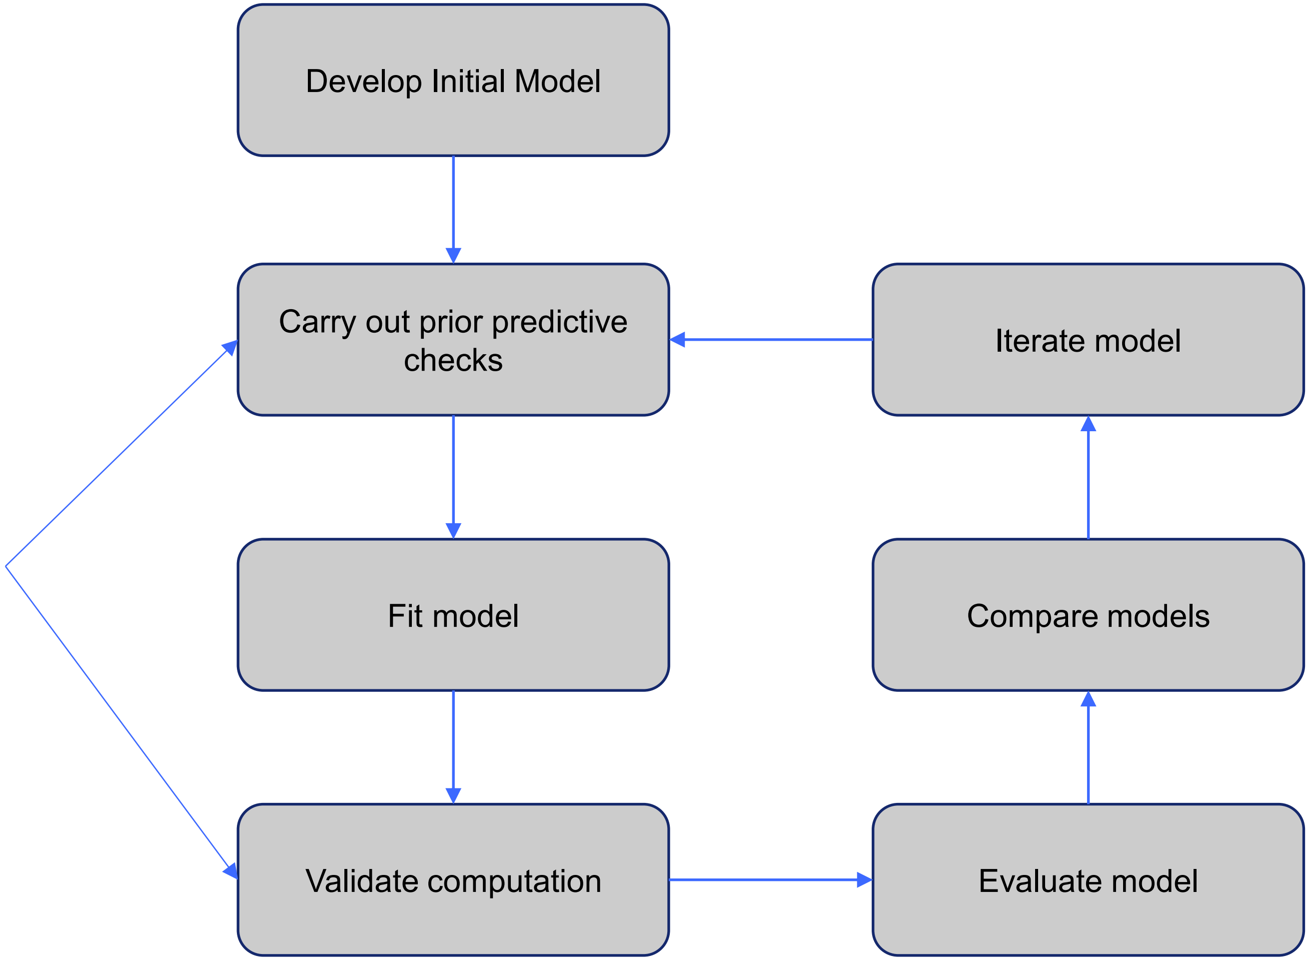

Case Study

Experience Study - Mortality and TPD Incidence Rates

Actuaries Institute.

Case Study – Mortality and TPD Experience

The objective of this case study is to show how, without much effort, we can build a single model that will:

- Fit a smooth age-based curve to mortality and TPD incidence rates
- Identify underlying changes in experience over time
- Implicitly allow for long reporting delays on TPD
- Provide estimates of uncertainty on the results

This following approach can be particularly valuable today for Retail TPD as experience is evolving fast. This provides a less subjective method that relies on the data, quantifies uncertainty and is easy to update.

Case Study – Experience Study

Simulated Data

Simulated claim experience from a hypothetical retail mortality and TPD portfolio across 2016-2025 for an Australian insurer with ~15% market share

Simulation

- Uses assumptions for mortality and TPD with some mortality improvement and TPD deterioration
- Reported claims simulated using TPD reporting delays
- Modelling uses a pricing/valuation basis different to the basis used to simulate claims
- Careful to make sure there is no data leakage

Case Study – Experience Study

Simulated Data

Claims experience by year

TPD actual claims are reported claims plus chain-ladder IBNR

	Mortality				TPD				
Year	Exposure	Expected Claims	Actual Claims	A/E	Exposure	Expected Claims	Reported Claims	Actual Claims	A/E
2016	355,959	763	718	94.1%	200,160	267	298	300	112.2%
2017	332,960	724	659	91.1%	192,016	258	342	344	133.7%
2018	305,958	670	578	86.3%	176,723	236	290	294	124.4%
2019	290,192	642	530	82.6%	169,743	228	260	266	116.5%
2020	260,927	579	509	87.9%	148,383	198	257	267	134.6%
2021	246,747	551	478	86.7%	141,028	190	261	281	147.5%
2022	232,891	523	420	80.4%	147,171	201	248	277	137.9%
2023	225,202	509	414	81.4%	146,338	201	248	305	151.9%
2024	201,244	453	326	71.9%	132,251	180	168	284	157.7%
2025	189,064	428	332	77.6%	128,785	177	49	305	172.5%
Total	**2,641,144**	**5,841**	**4,964**	**85.0%**	**1,582,598**	**2,137**	**2,421**	**2,923**	**136.8%**

Claims experience by age band

TPD actual claims are reported claims plus chain-ladder IBNR

	Mortality				TPD				
Age Band	Exposure	Expected Claims	Actual Claims	A/E	Exposure	Expected Claims	Reported Claims	Actual Claims	A/E
20-24	14,103	5	8	158.7%	20,023	4	7	7	201.4%
25-29	64,036	27	28	103.9%	79,408	15	16	17	113.7%
30-34	153,901	84	75	89.3%	153,638	46	45	56	123.8%
35-39	276,911	210	172	81.8%	233,841	111	103	123	110.8%
40-44	399,626	441	383	86.8%	280,107	205	221	279	136.2%
45-49	495,841	805	655	81.3%	274,238	313	299	353	112.7%
50-54	510,641	1,205	1,048	87.0%	242,164	434	456	564	129.8%
55-59	449,807	1,607	1,341	83.4%	188,060	526	600	720	136.7%
60-64	276,278	1,456	1,254	86.1%	111,119	483	674	804	166.6%
Total	**2,641,144**	**5,841**	**4,964**	**85.0%**	**1,582,598**	**2,137**	**2,421**	**2,923**	**136.8%**

Case Study – Experience Study

Overview of Experience

- Simulated claim experience from a hypothetical retail mortality and TPD portfolio across 2016-2025 for an Australian insurer with ~15% market share

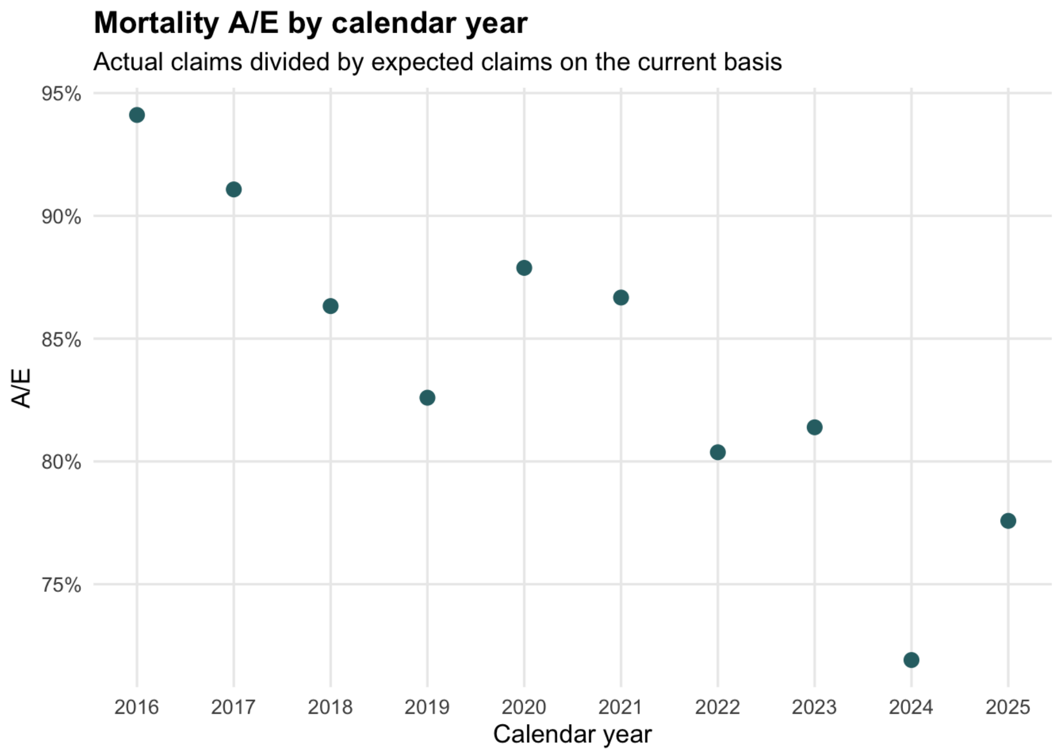

Mortality is gradually improving

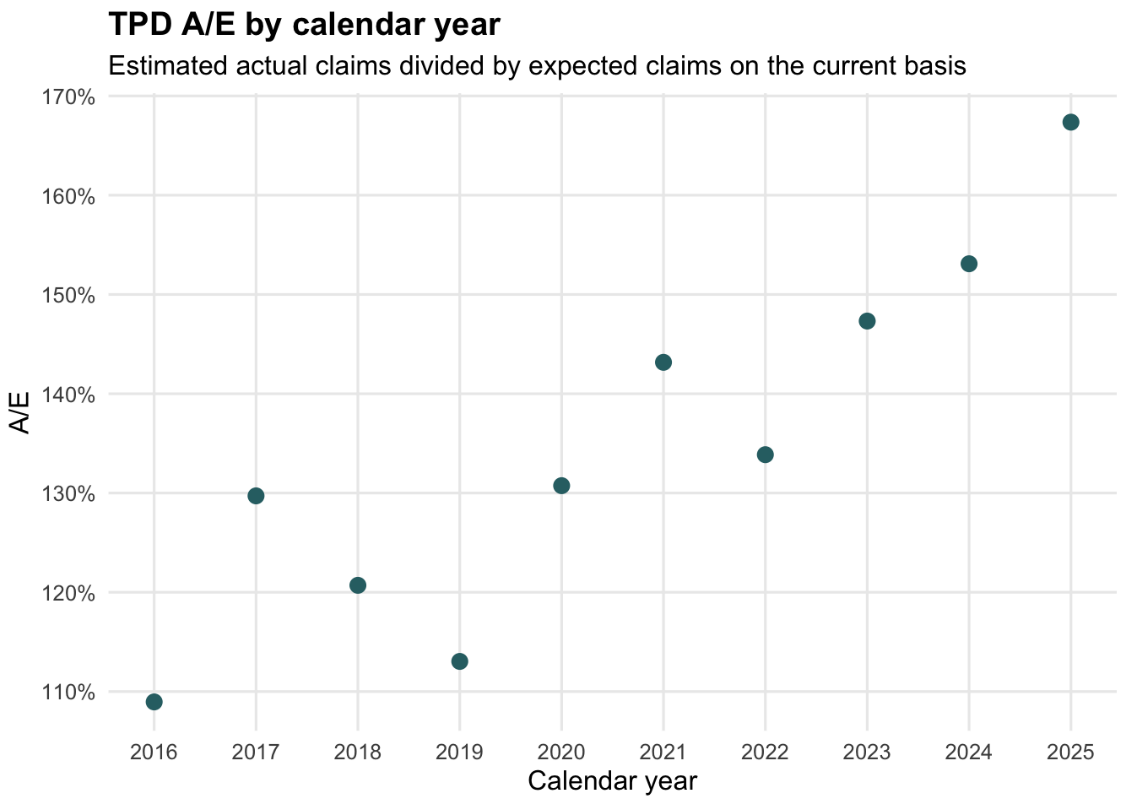

TPD is rapidly deteriorating

Case Study cont.

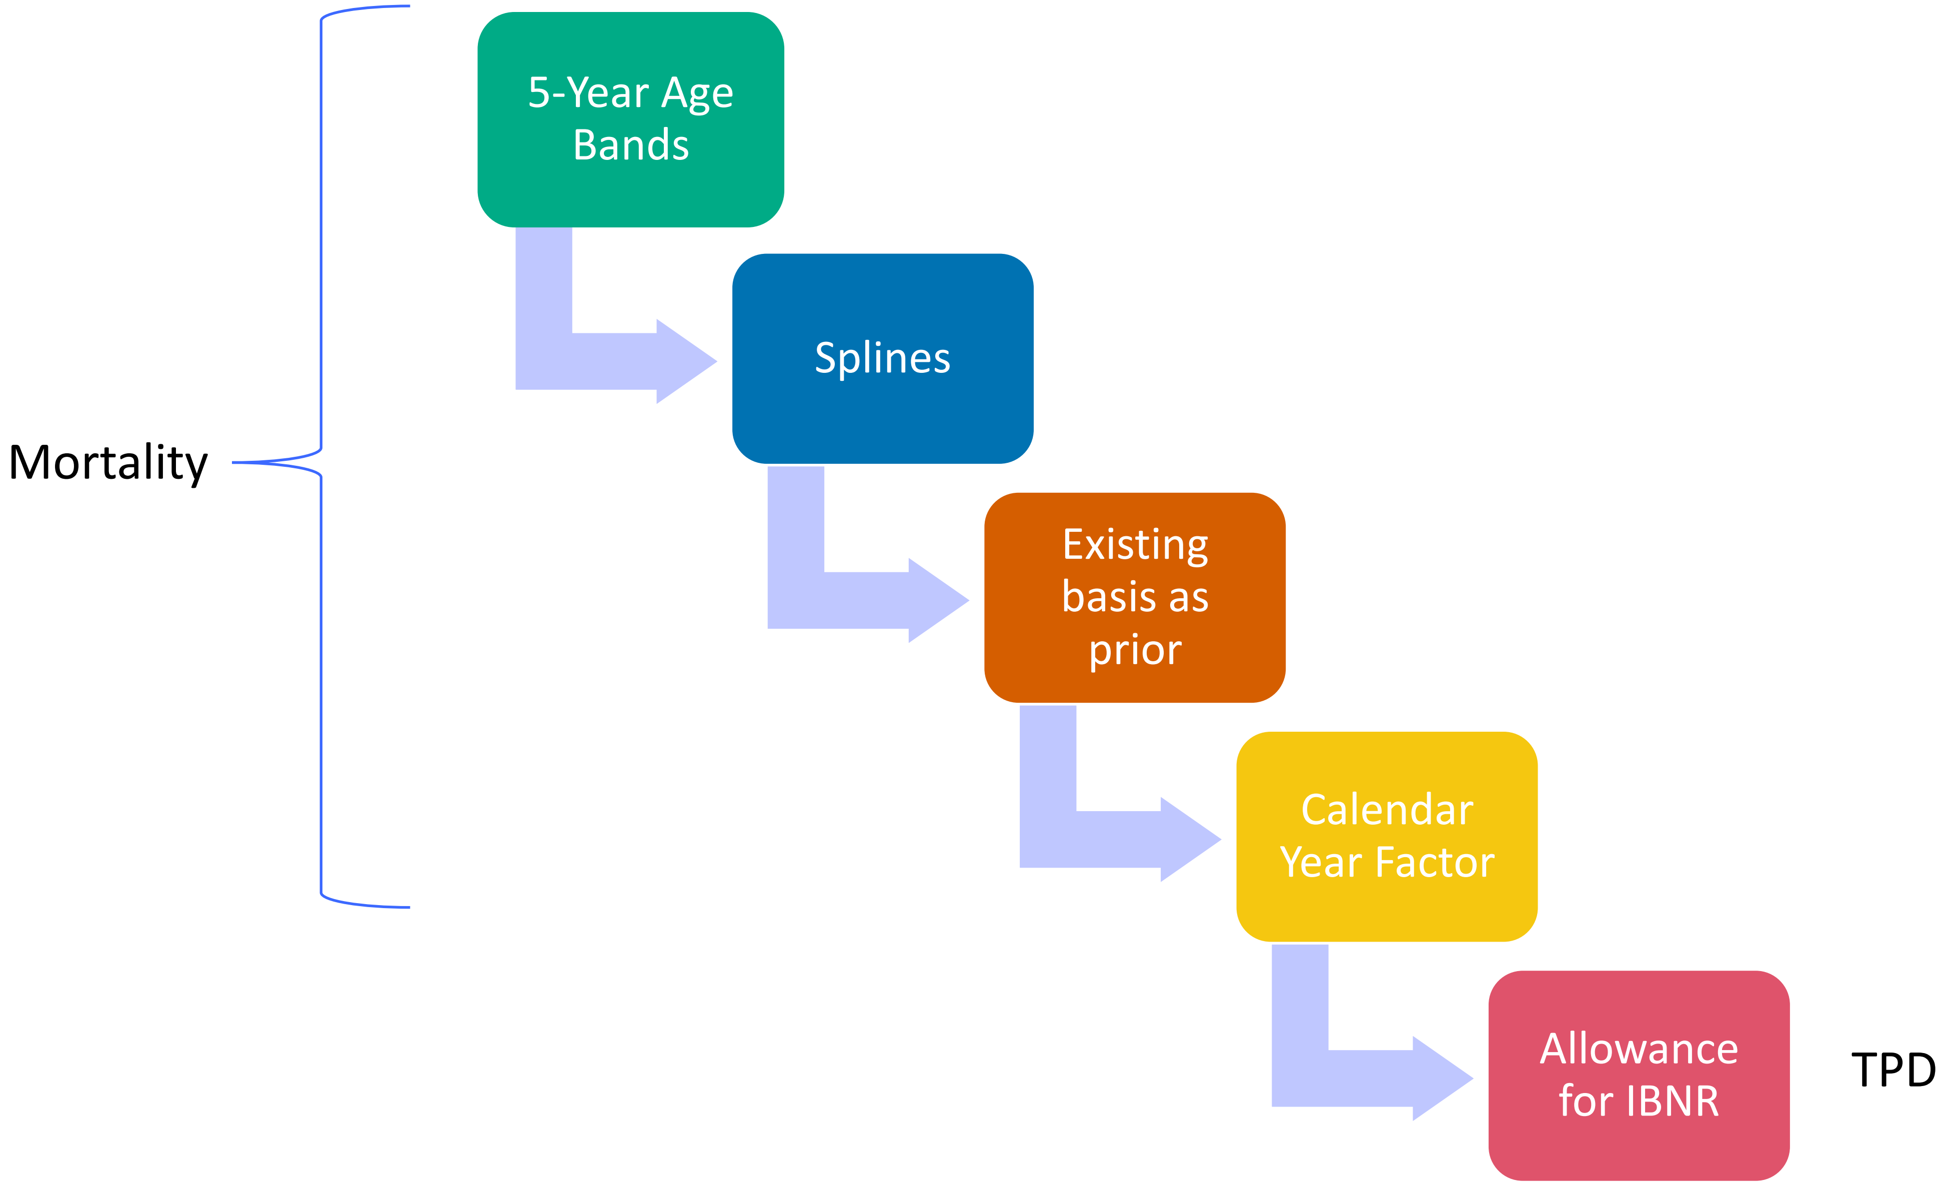

Case Study continued.

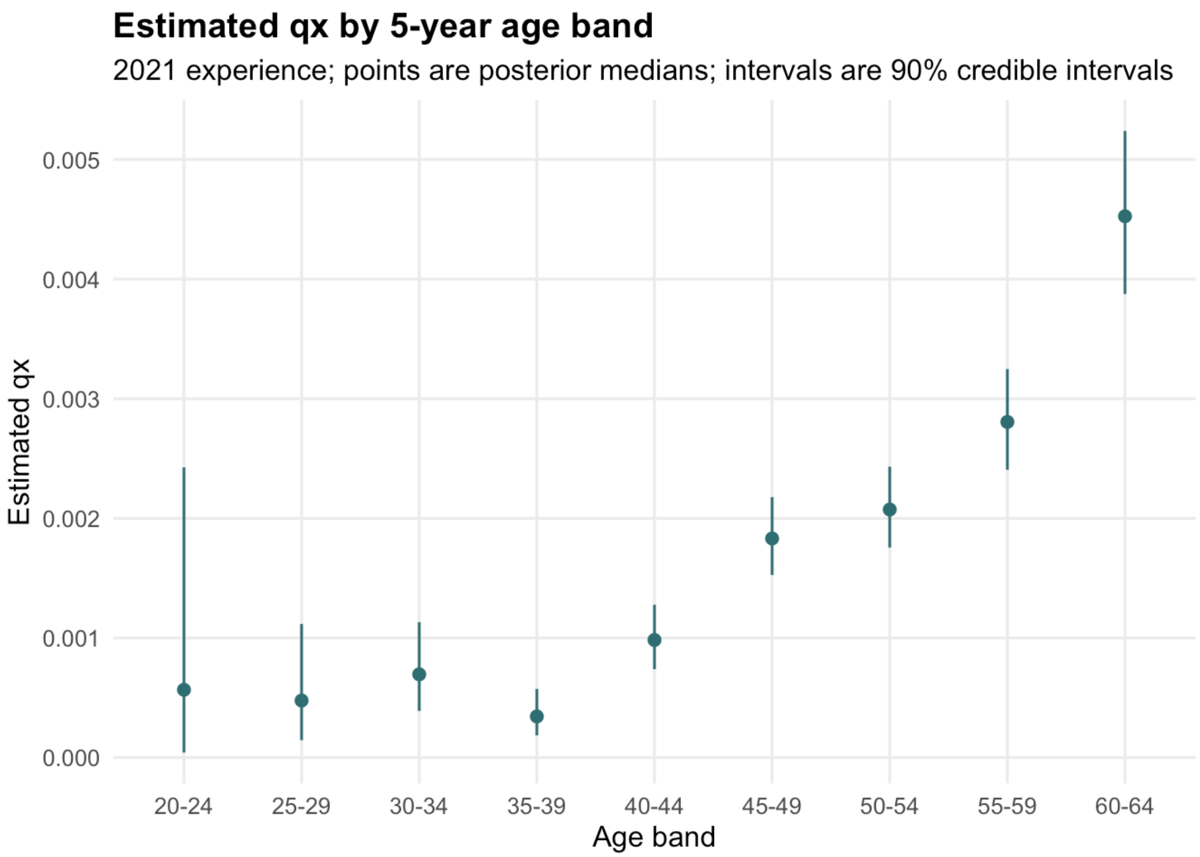

Model:

$q_x \sim \text{beta}(1, 100)$

$\text{deaths}_x \sim \text{binomial}(\text{exposure}_x, q_x)$

Case Study continued.

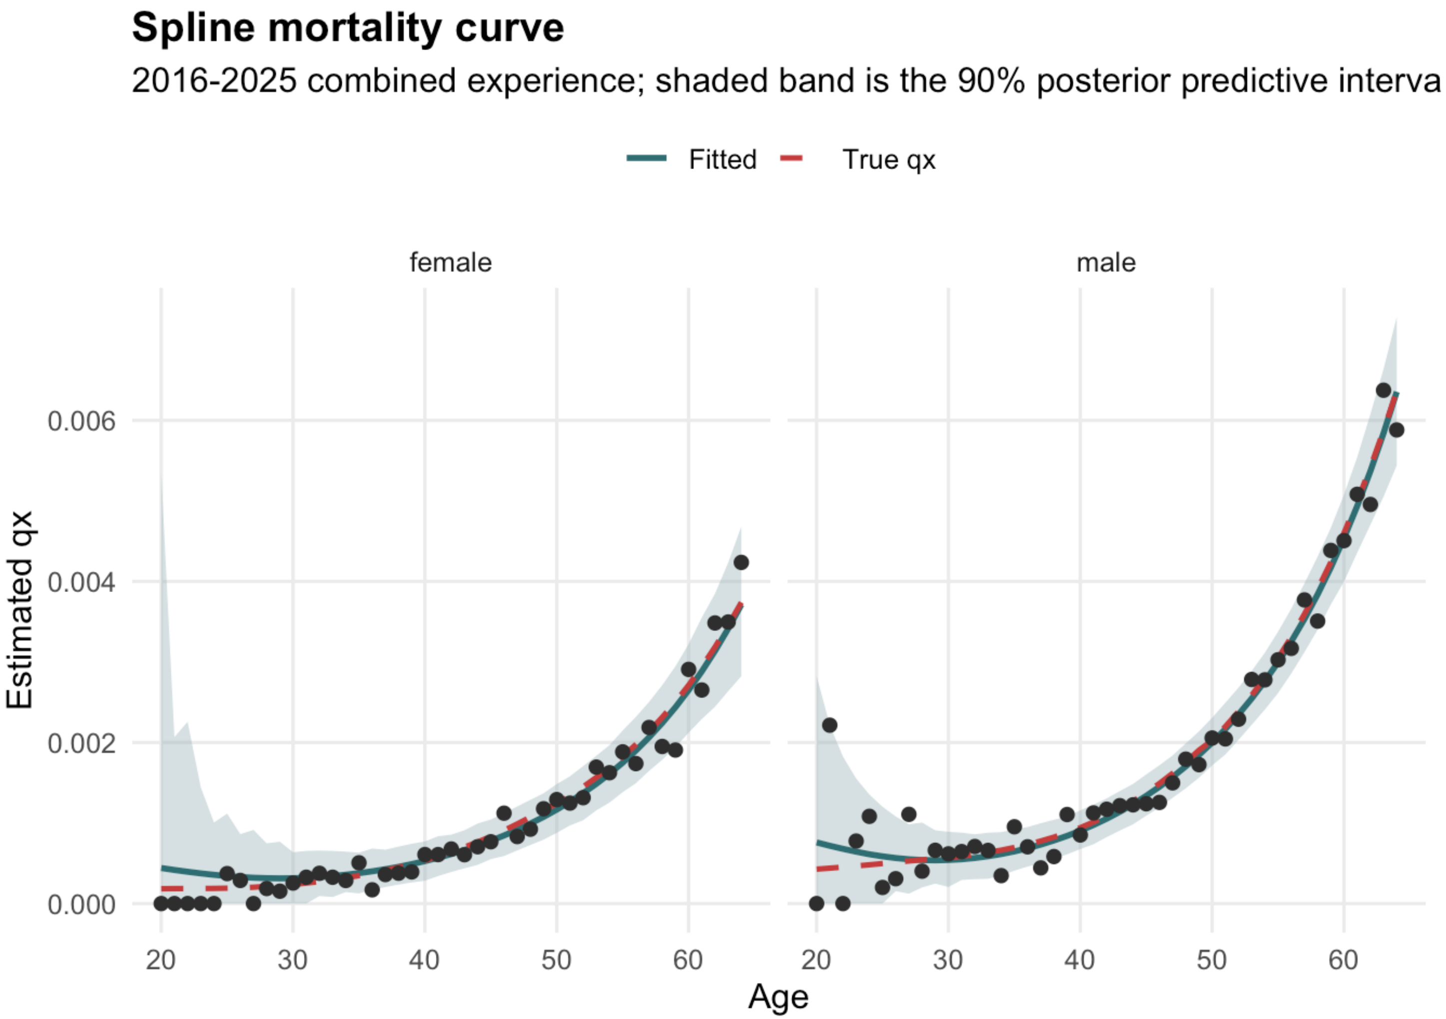

Priors:

beta ~ normal(0, 10)

Model:

eta = X * beta

q_x = inv_logit(eta_x)

$deaths_x$ ~ binomial($exposure_x$, q_x)

Where

X is a N x K design matrix for the splines

beta[K] is the vector of spline basis coefficients

Use a logit scale as q_x is a probability between 0 and 1.

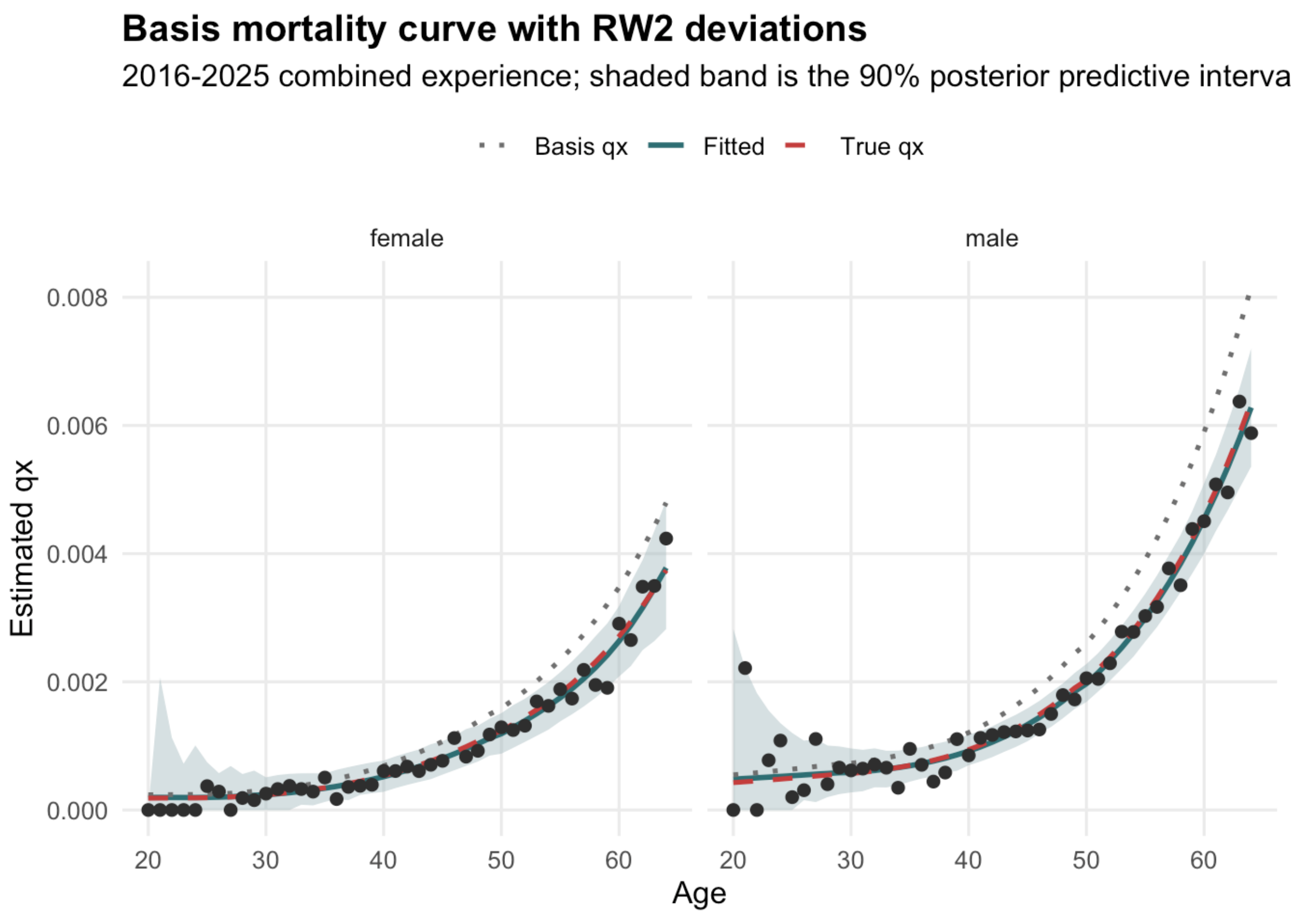

Priors:

overall_level ~ normal(0, 0.5)

age_adj_sigma ~ normal(0, 0.2)

age_adj[1] ~ normal(0, 0.25)

age_adj[2] ~ normal(0, 0.25)

age_adj[x] ~ normal(2 * age_adj[x-1] – age_adj[x-2], age_adj_sigma)

Model:

qx[x] = qx_basis[x] * exp(overall_level + **age_adj[x] – mean(age_adj)**)

deaths ~ binomial(exposure, qx)

Where

qx_basis equals current pricing/valuation/industry basis

Basis mortality curve with age and calendar RW2 deviations

2016-2025 combined experience; shaded band is the exposure-weighted 90% poste

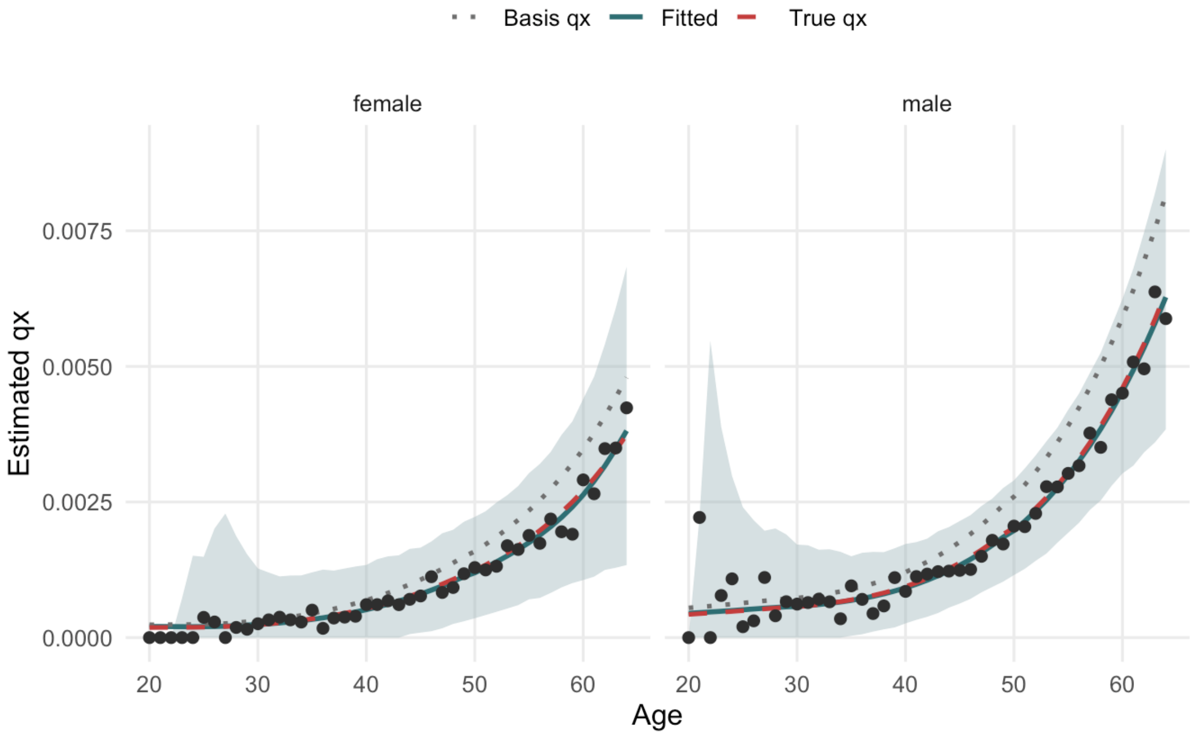

Calendar year qx factor vs basis

Fitted expected claims divided by expected claims on the current basis

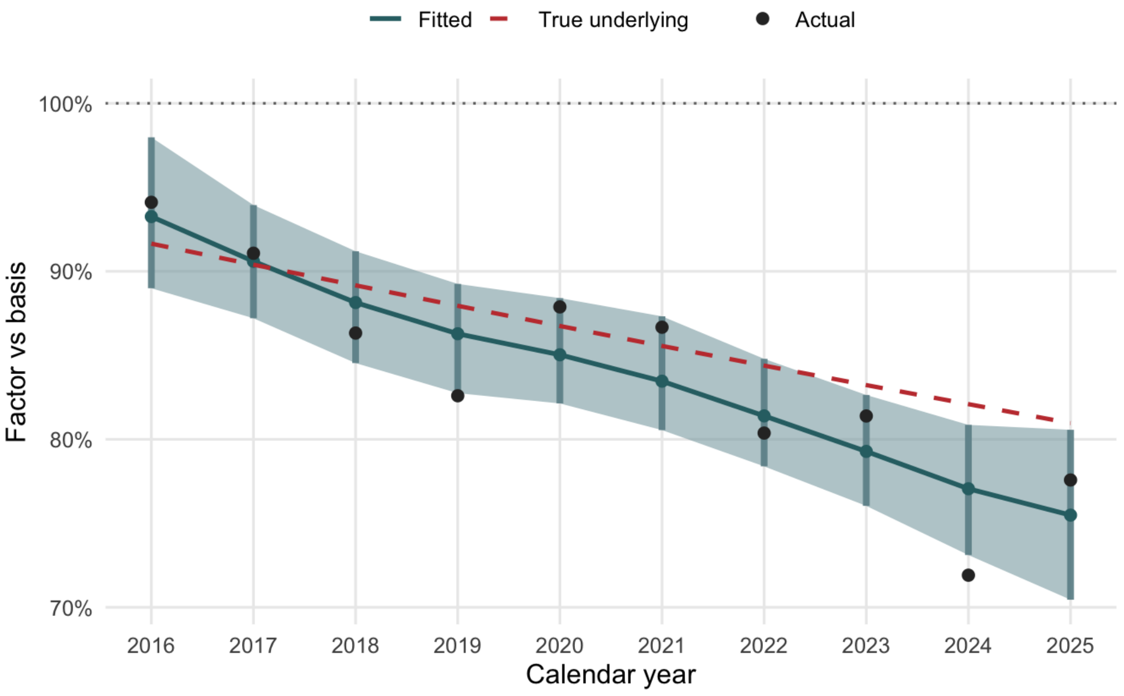

Priors:

overall_level ~ normal(0, 0.5)

age_adj_sigma ~ normal(0, 0.2)

age_adj[1] ~ normal(0, 0.25);   age_adj[2] ~ normal(0, 0.25)

age_adj[x] ~ normal(2 * age_adj[x-1] – age_adj[x-2], age_adj_sigma)

keeps adjustment to given age curve smooth

cal_adj_sigma ~ normal(0, 0.2)

cal_adj[1] ~ normal(0, 0.25);   cal_adj[2] ~ normal(0, 0.25)

cal_adj[x] ~ normal(2 * cal_adj[x-1] – cal_adj[x-2], cal_adj_sigma)

keeps calendar effects smooth

Model:

qx[x] = qx_basis[x] * exp(overall_level + **age_adj[x] – mean(age_adj)** + **cal_adj[x] – mean(cal_adj)**)

deaths ~ binomial(exposure, qx)

Where

qx_basis equals current pricing/valuation/industry basis

Case Study – TPD Experience

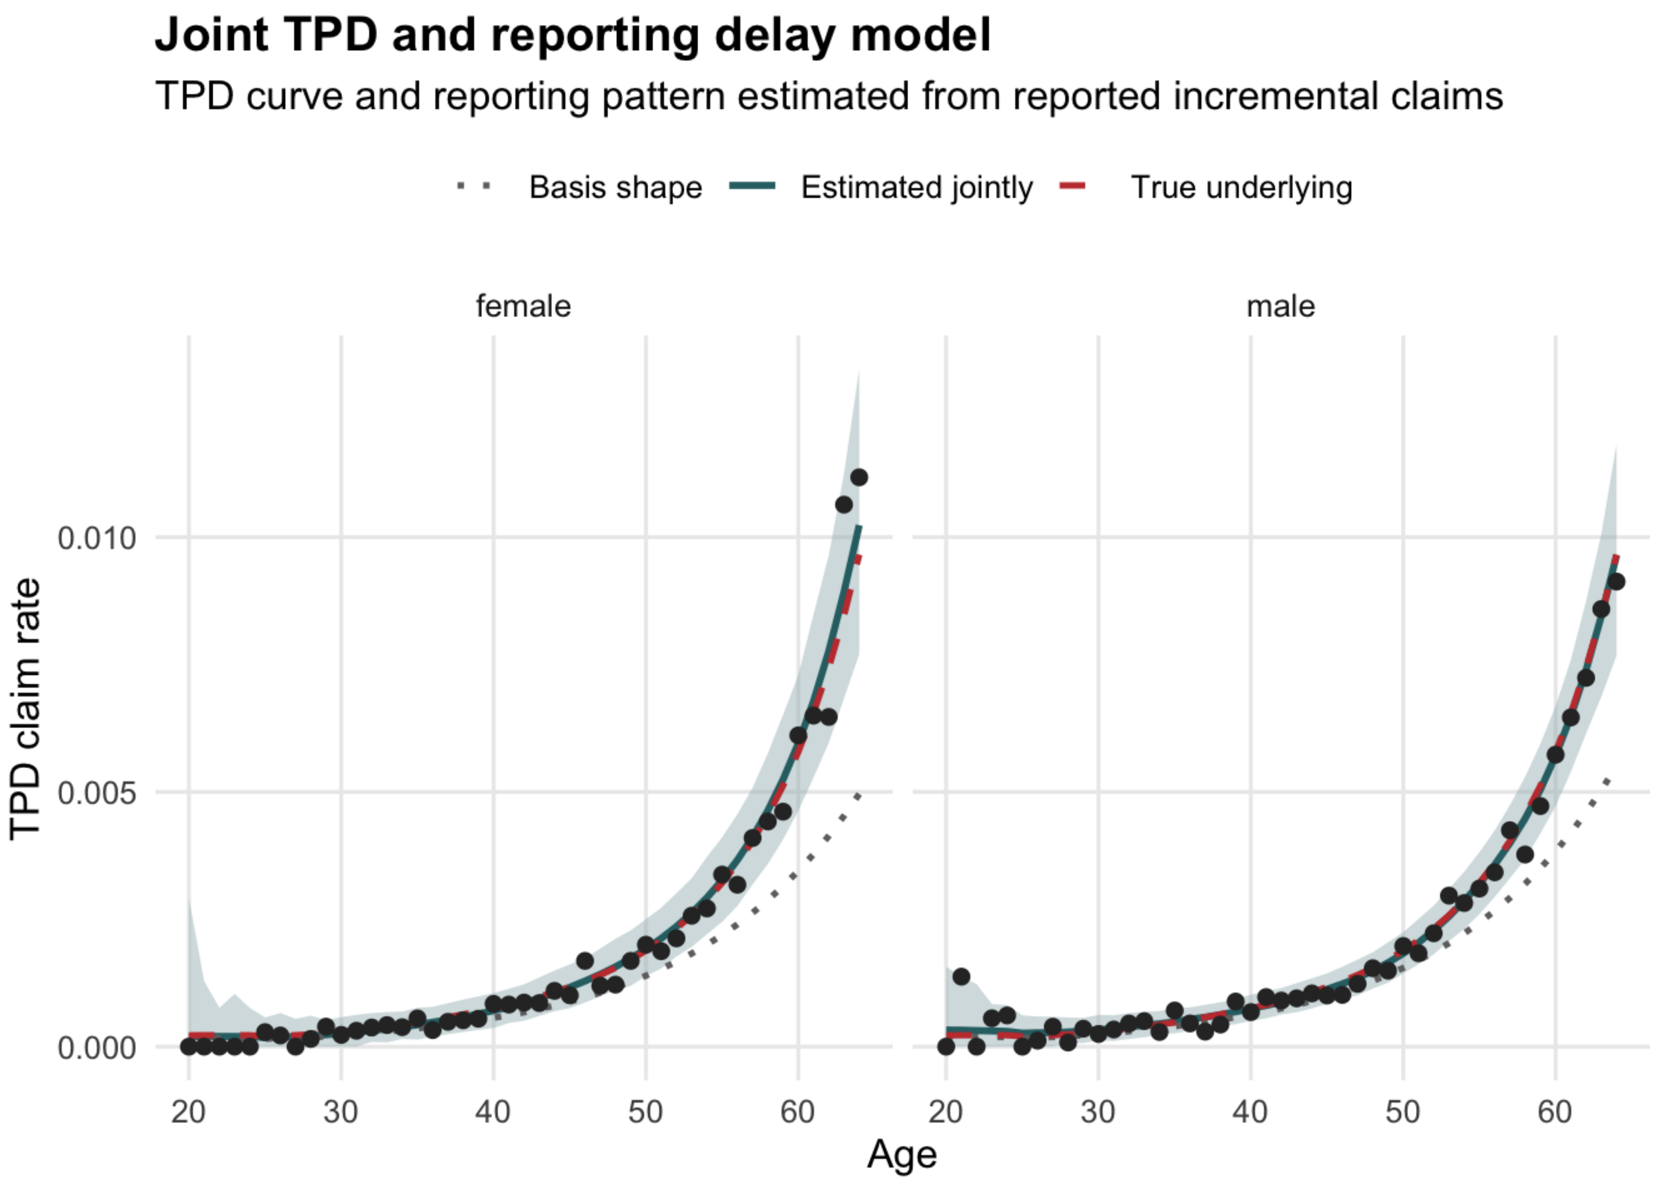

Case Study – TPD Experience

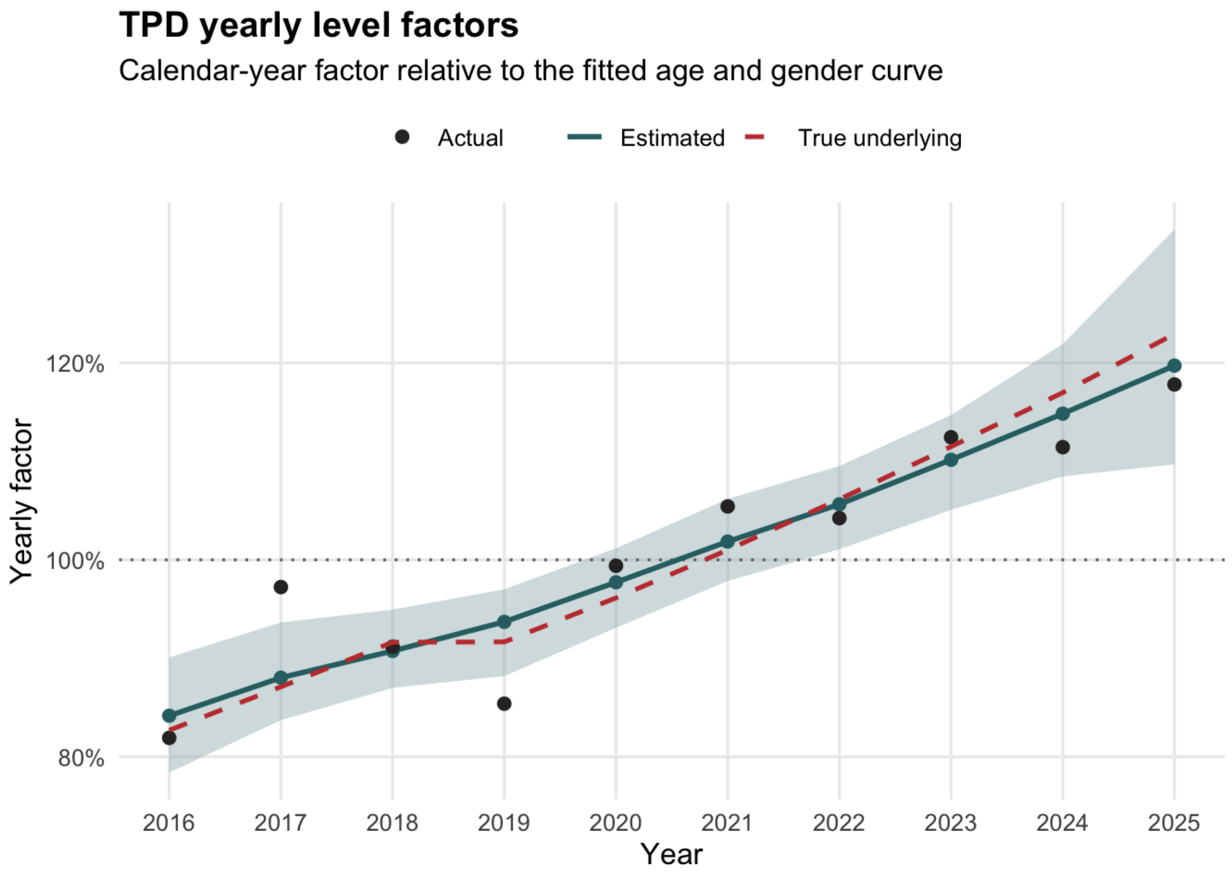

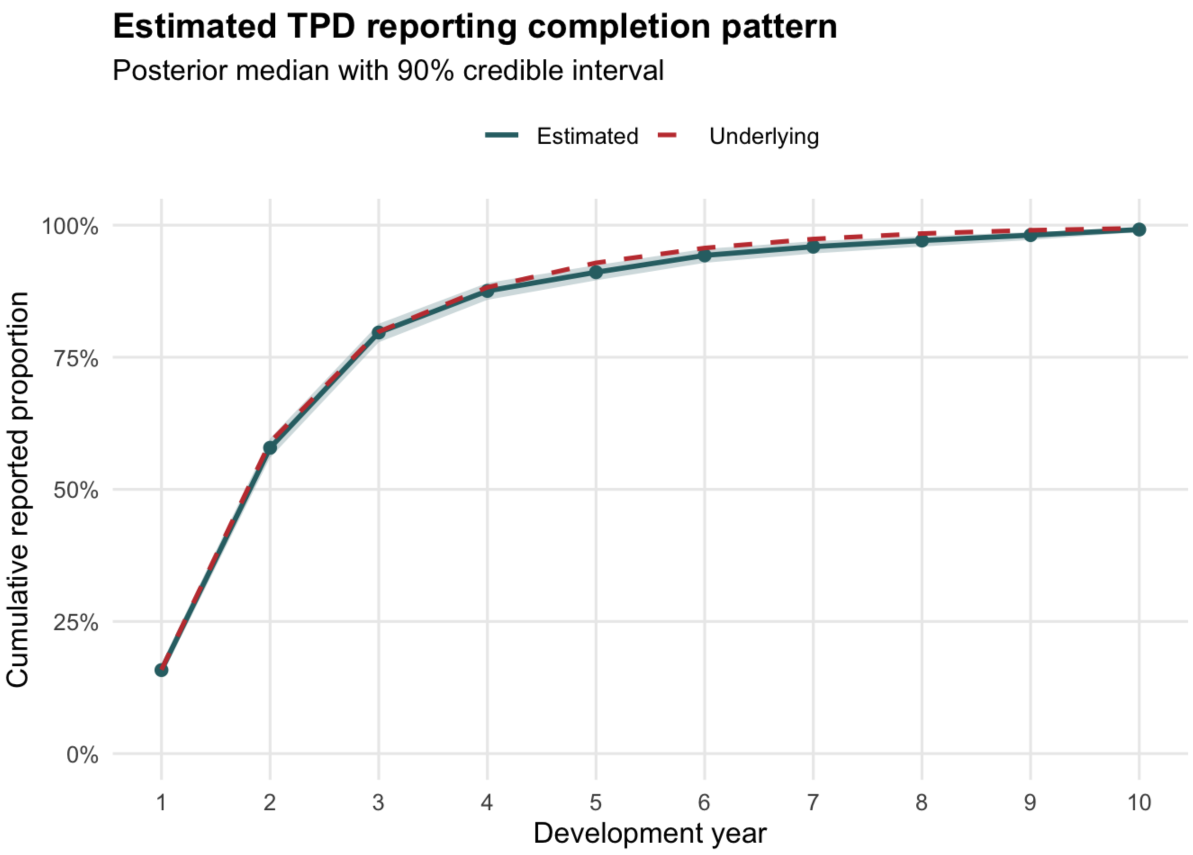

Case Study – TPD Experience Model

```
Priors:
  overall_level ~ normal(0, 0.5)
  age_adj_sigma ~ normal(0, 0.2)
  age_adj[1] ~ normal(0, 0.25);    age_adj[2] ~ normal(0, 0.25)
  age_adj[x]  ~ normal(2 * age_adj[x-1] – age_adj[x-2], age_adj_sigma)
  cal_adj_sigma ~ normal(0, 0.2)
  cal_adj[1] ~ normal(0, 0.25);    cal_adj[2] ~ normal(0, 0.25)
  cal_adj[x] ~ normal(2 * cal_adj[x-1] – cal_adj[x-2], cal_adj_sigma)
  reporting_prob ~ dirichlet(reporting_prior)

Model:
  qx[x] = tpd_basis[x] * exp(overall_level + age_adj[x] – mean(age_adj) + cal_adj[x] – mean(cal_adj))
  tpd_claims ~ poisson(exposure * qx * reporting_prob)

Where
  qx_basis equals current pricing/valuation/industry basis
```

keeps adjustment to given age curve smooth

keeps calendar effects smooth

Fitted incidence curves by year

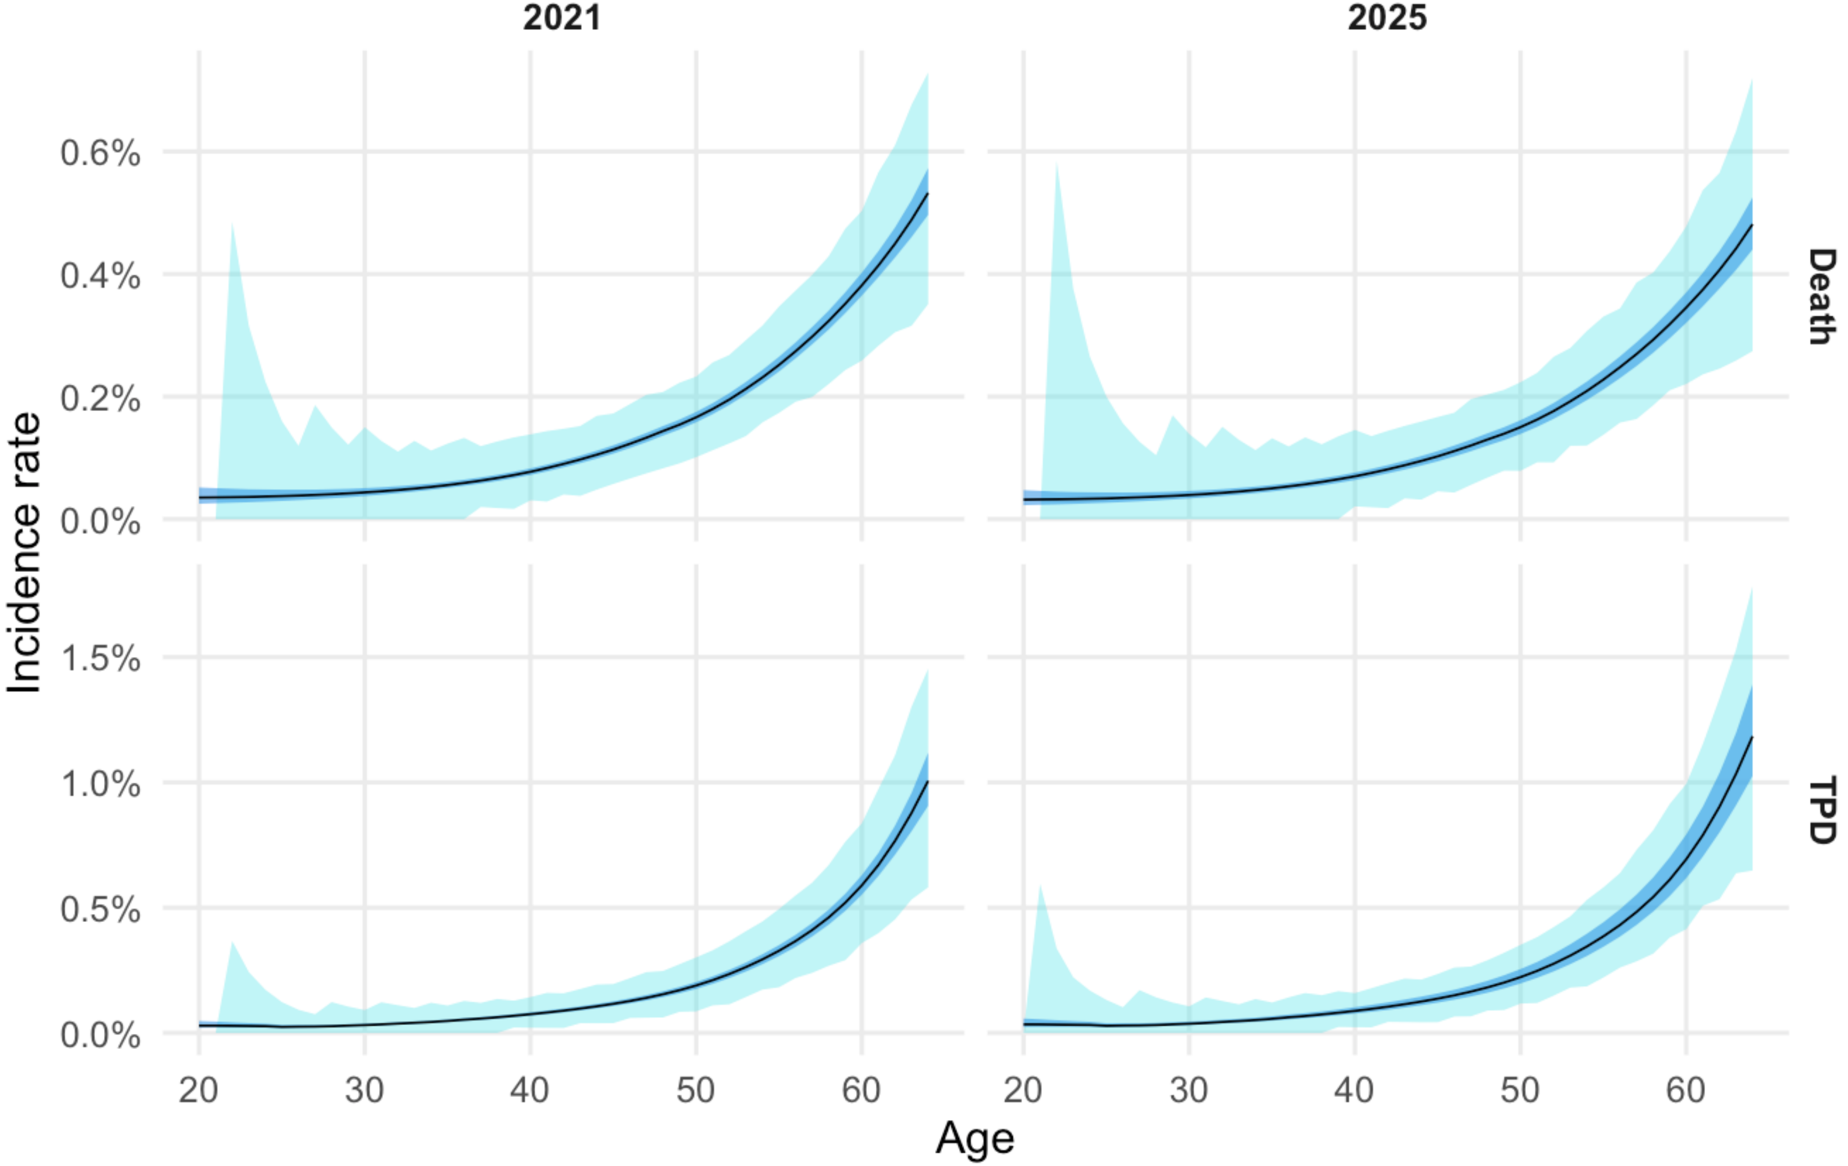

Dark ribbon: 90% credible interval; Light ribbon: 90% posterior predictive interval

Total expected claims relative to basis - 2025

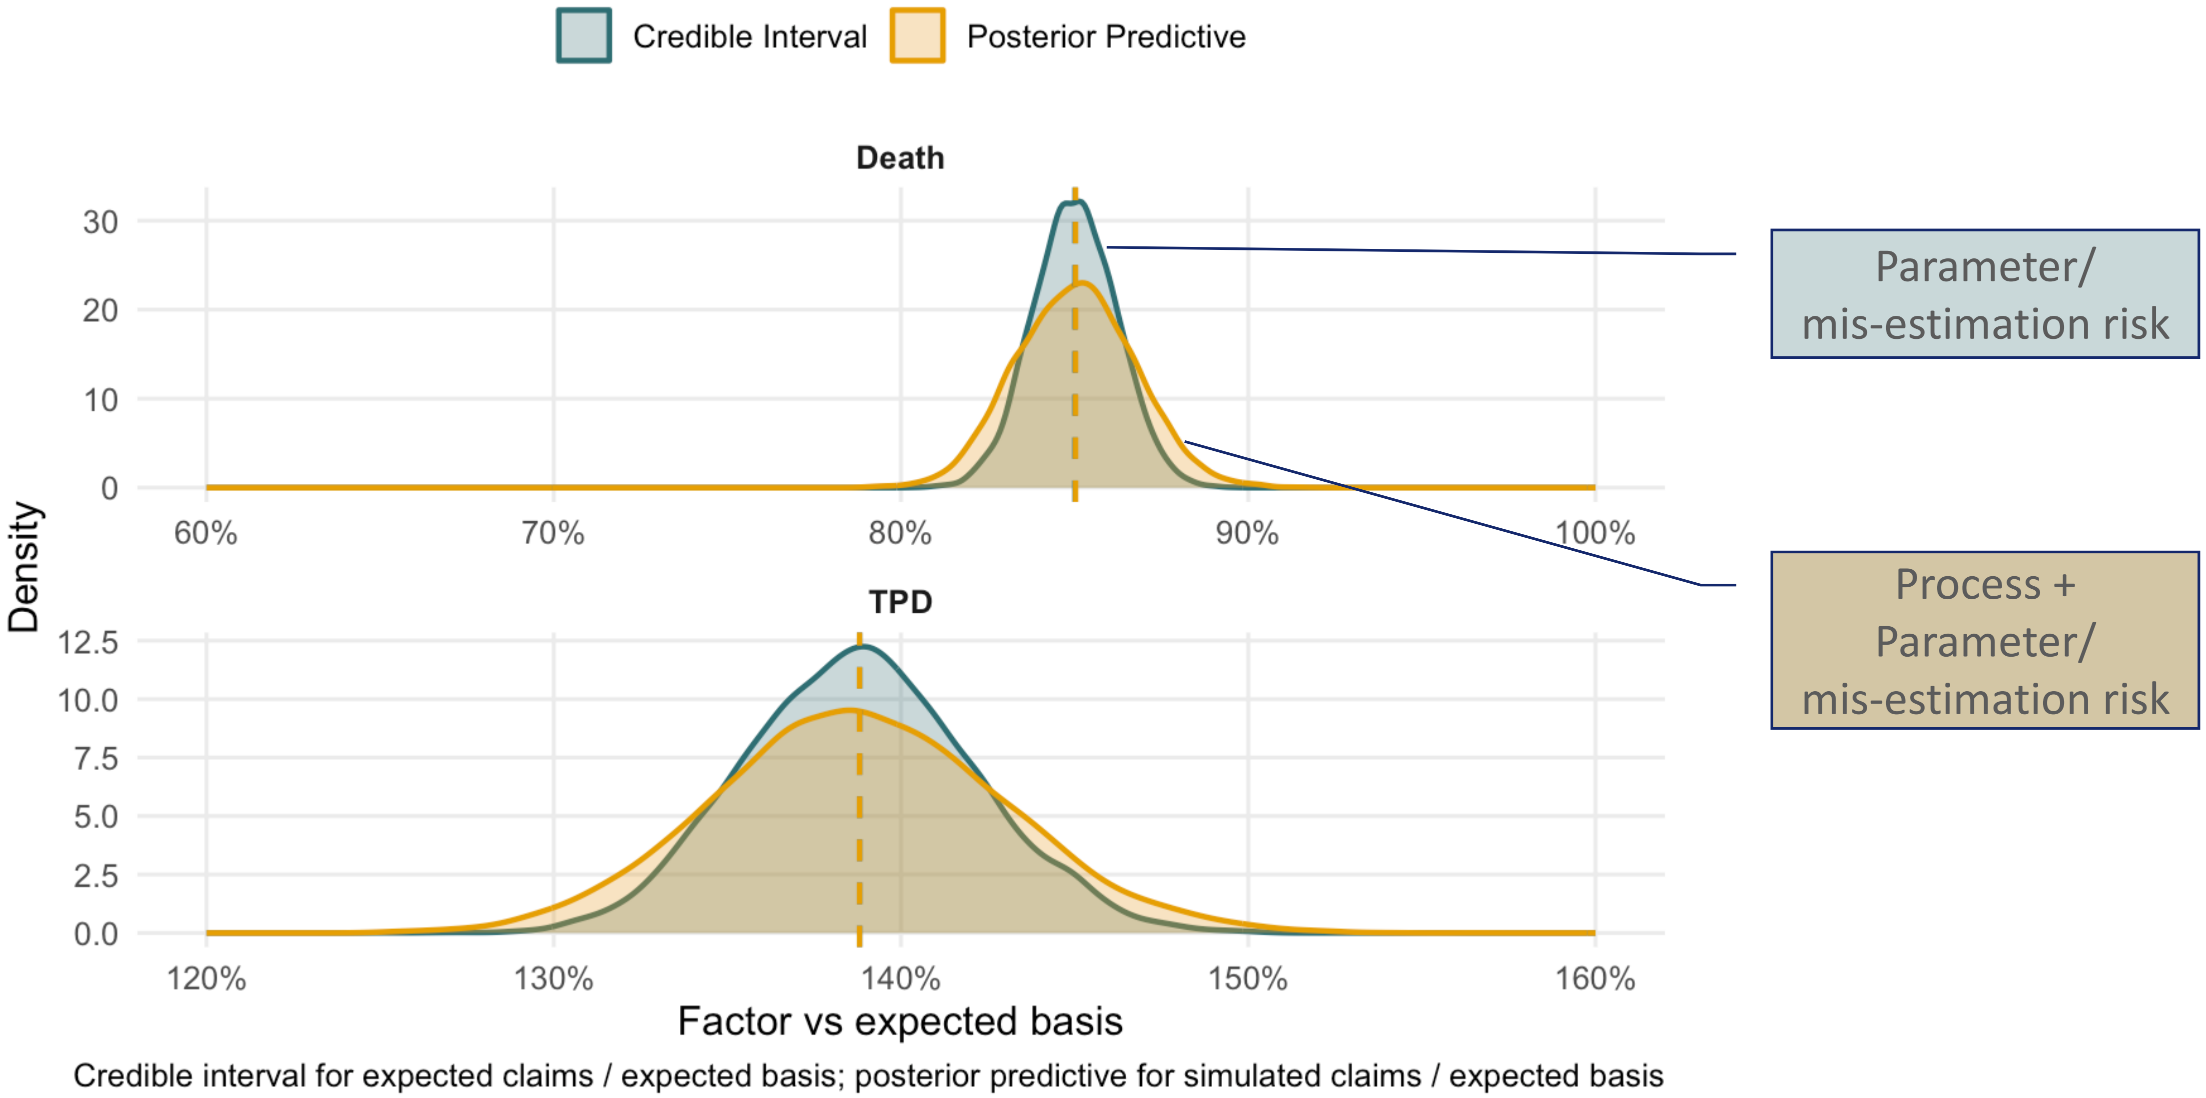

Case Study - Summary

This case study shows:

- How we can justify base changes more rigorously
- Use emerging experience without over-reacting to it
- Separate true experience trends from reporting delays
- Quantify uncertainty for valuation pricing and risk margins
- Make actuarial judgement more transparent

The value of the Bayesian approach is that it gives actuaries a disciplined way to combine current basis, emerging experience and judgement while preserving uncertainty for the decisions that depend on it.

This approach can be particularly valuable today for Retail TPD when experience is evolving fast and you want a less subjective method that relies on the data and quantifies uncertainty.

Similar approaches can be applied to:

- Group Pricing
- Mortality Improvements

Case Study - Extensions

Group Pricing

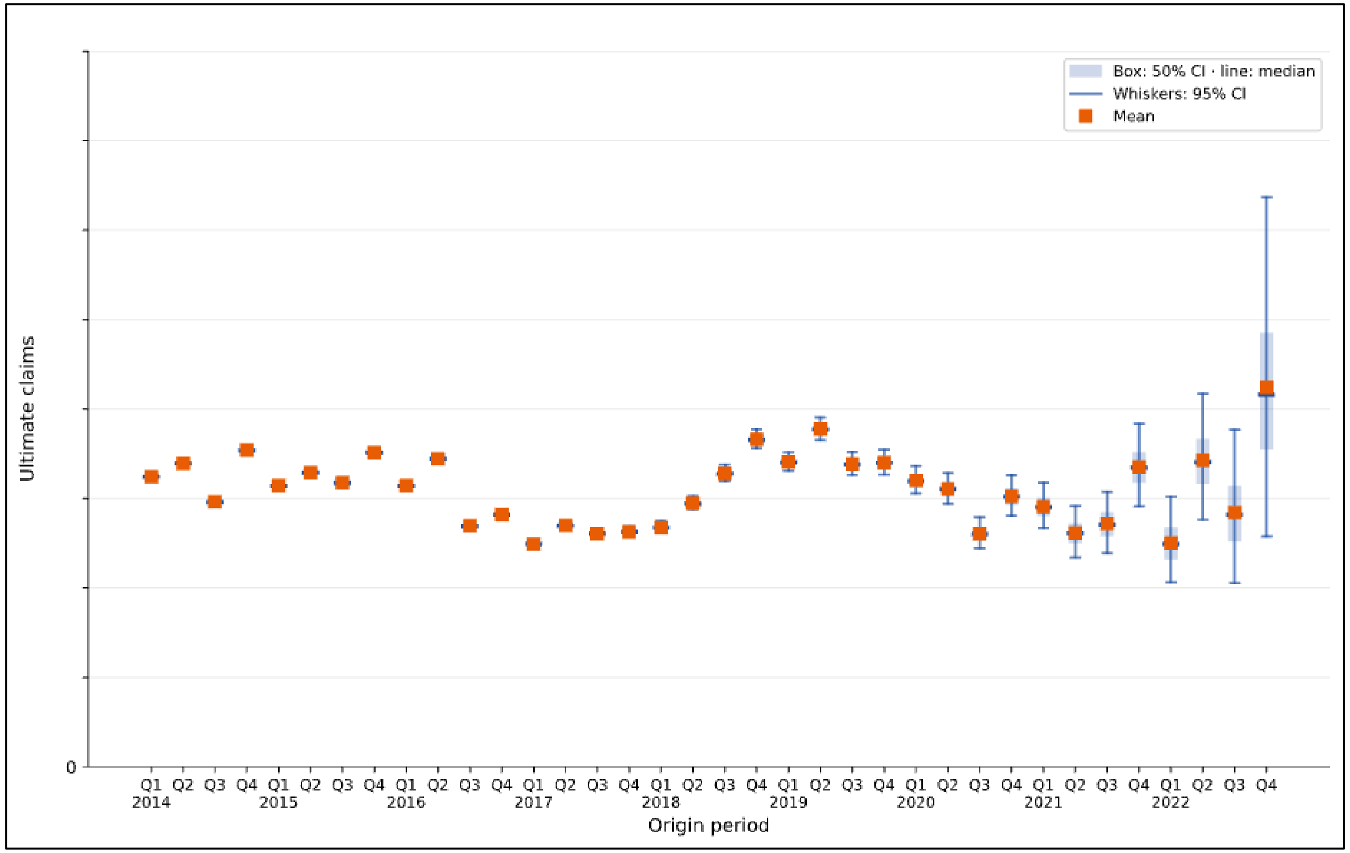

Mortality Improvements

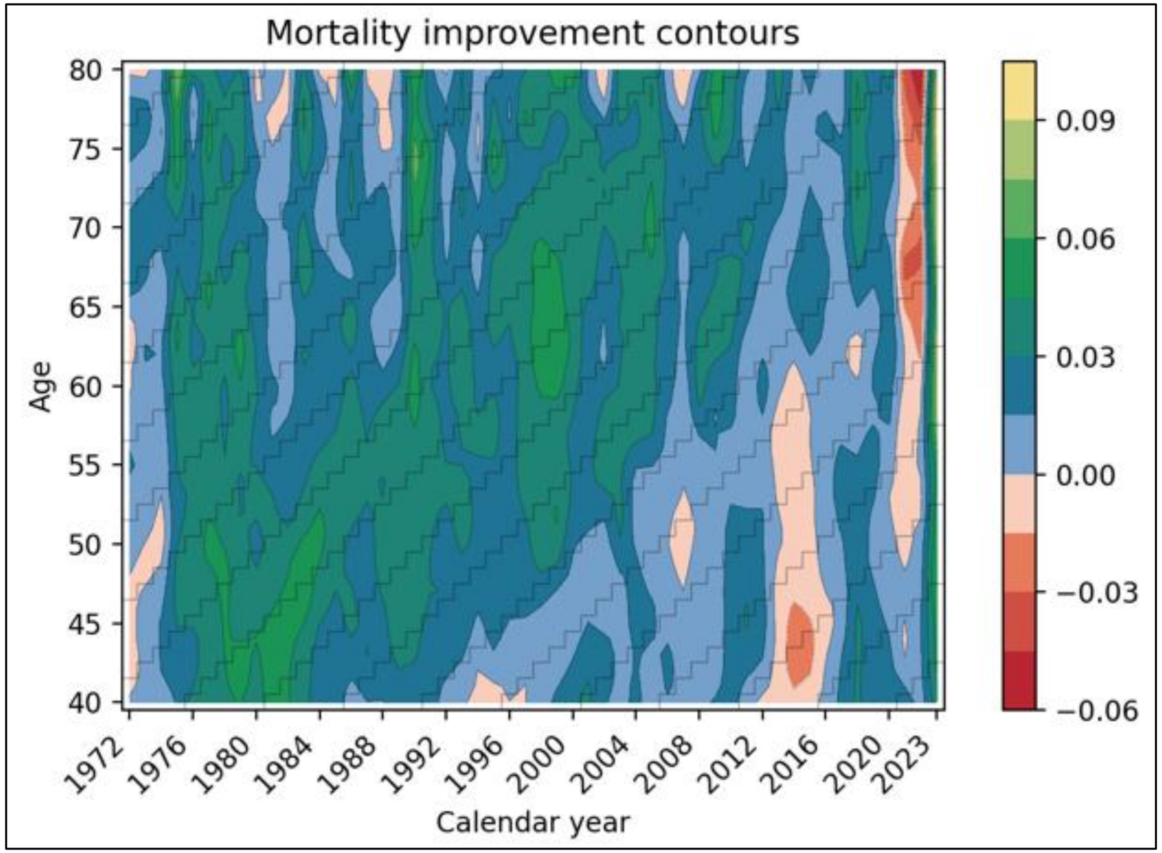

Areas of applications

Experience sharing mechanism pricing

- Loss distributions are a key input into valuing experience sharing mechanisms

Anti-selection

- Can capture anti-selective effect as a latent variable on segments of a portfolio

Assessment of risk margins

- Estimate risk margins particularly for IFRS17 and LAGIC margins related to mis-estimation and random stresses.

Experience studies

- Incorporating the existing valuation/pricing basis into the prior can provide a solid statistical foundation to justify basis adjustments.

Corporate group pricing

- A Bayesian credibility model can be used to calculate premium that would model the correct loss distribution and would recalibrate assumptions with each new scheme added to the portfolio.

Mortality modelling

- A Bayesian model can fit historical factors and long-term improvement factors simultaneously using a single model to reduce information loss. The second-order random walk is an alternative to cubic splines and reduces the need for arbitrary judgement.

Software requirements

Bayesian models are typically not analytically tractable meaning that estimation methods are required to estimate their posterior distributions.

The most popular method used to estimate a Bayesian model's posterior distribution is MCMC. MCMC is a complicated simulation-based method as a result this method is best fit using dedicated software.

Common software used to run MCMC are:

- Stan – this can be run from R, Python, Julia etc.
- PyMC – a python specific library
- Pyro – a python specific library
- Turing – Juilia specific library

We've found Stan to be the easiest to use as the syntax allows users to specify the model close to it mathematical form however it is not the best are every Bayesian modelling problem.