Measuring loss reserving uncertainty with machine learning models

Gráinne McGuire

Acknowledgements

- Joint work with Greg Taylor, University of New South Wales, Australia
- Greg's acknowledgements
  - financial support under Australian Research Council's Linkage Projects funding scheme (project number LP130100723)
  - Discussion with colleagues: Benjamin Avanzi, Bernard Wong, David Yu
- Gráinne's acknowledgements:
  - Support from Taylor Fry for computational work

Overview

- Introduction
- Claims reserving
- Uncertainty
- Lassoing the model set
- Bootstrapping
- Results
- Conclusion

Reference material

- Paper:
  - Model error (or Ambiguity) and its estimation with particular application to loss reserving
  - https://doi.org/10.3390/risks11110185, Risks 2023, 11(11), 185

- Tutorial example, including full R code:
  - Model error via regularised regression - CAS monograph data
  - https://grainnemcguire.github.io/post/2023-05-04-model-error-example/

Introduction

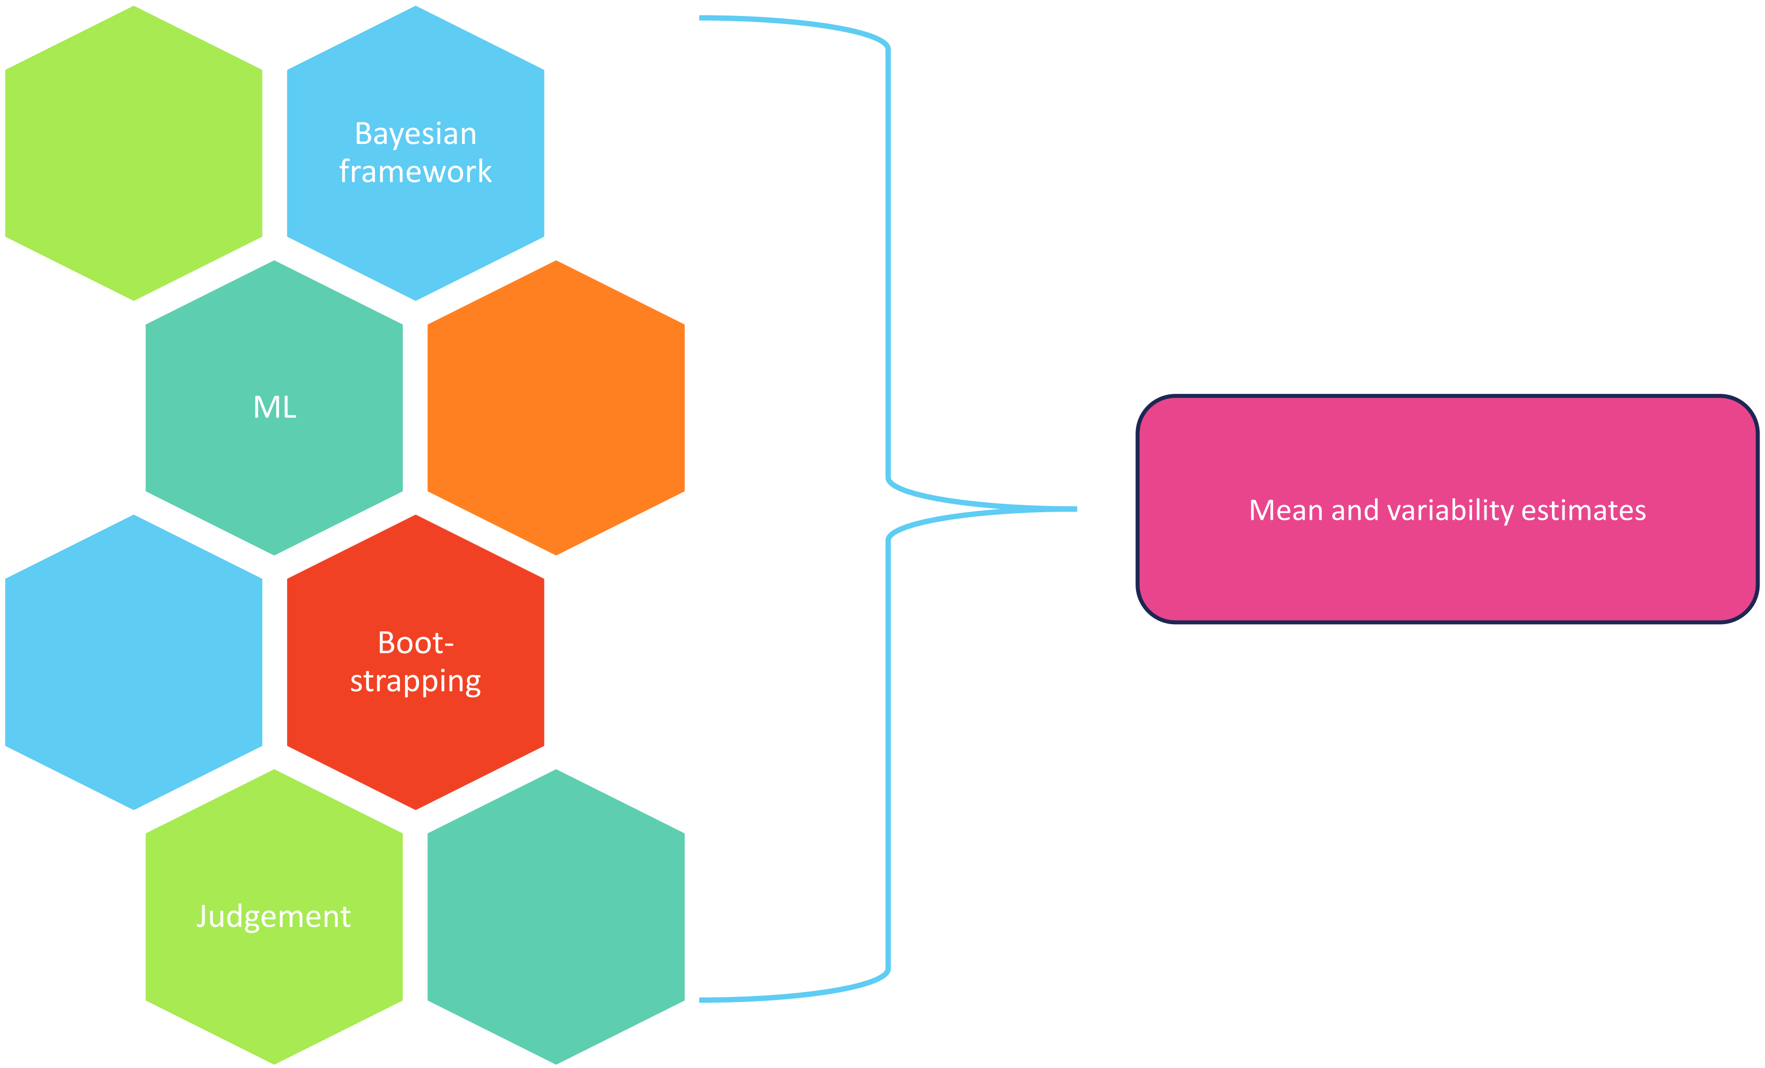
Bayesian framework
ML
Boot-strapping
Judgement
Mean and variability estimates

Claims reserving

Claims reserving problem

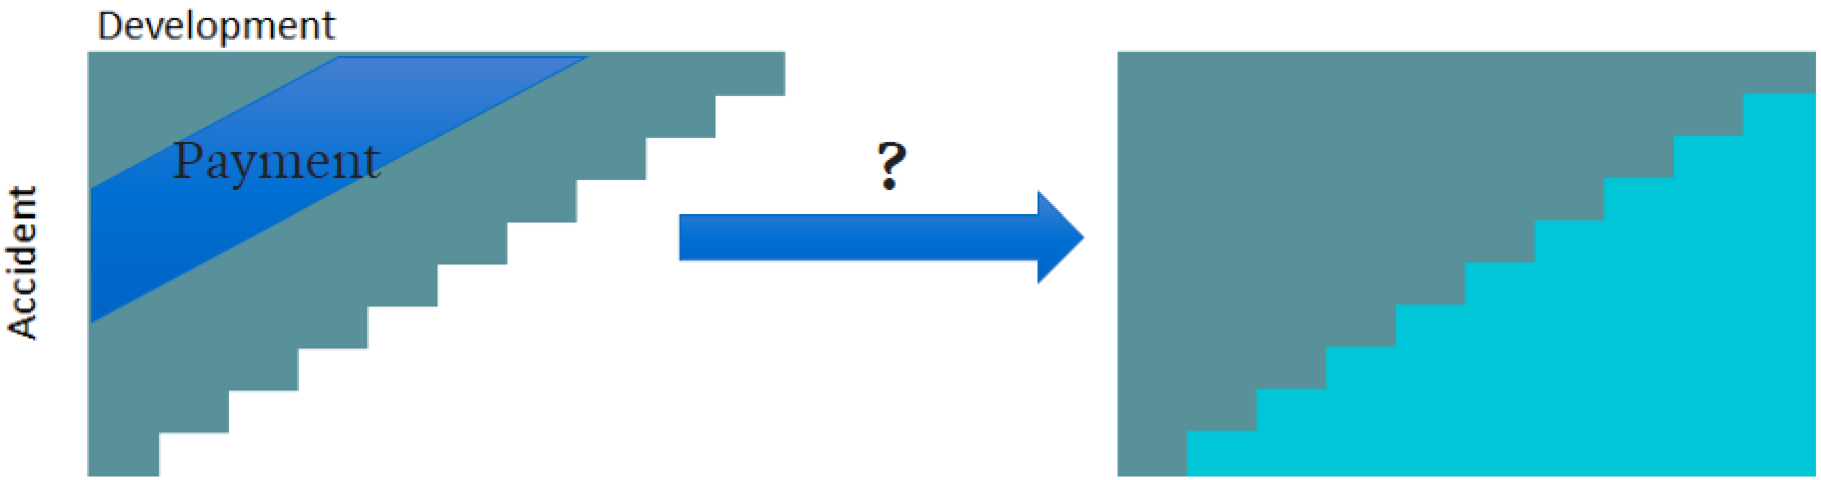

Long form

acc	dev	cal	pmt
1	1	1	0
1	2	2	100
1	3	3	200
.	.	.	.
.	.	.	.
9	2	10	25
10	1	10	50

Uncertainty

Forecast error = Observation − Forecast

Divide and conquer

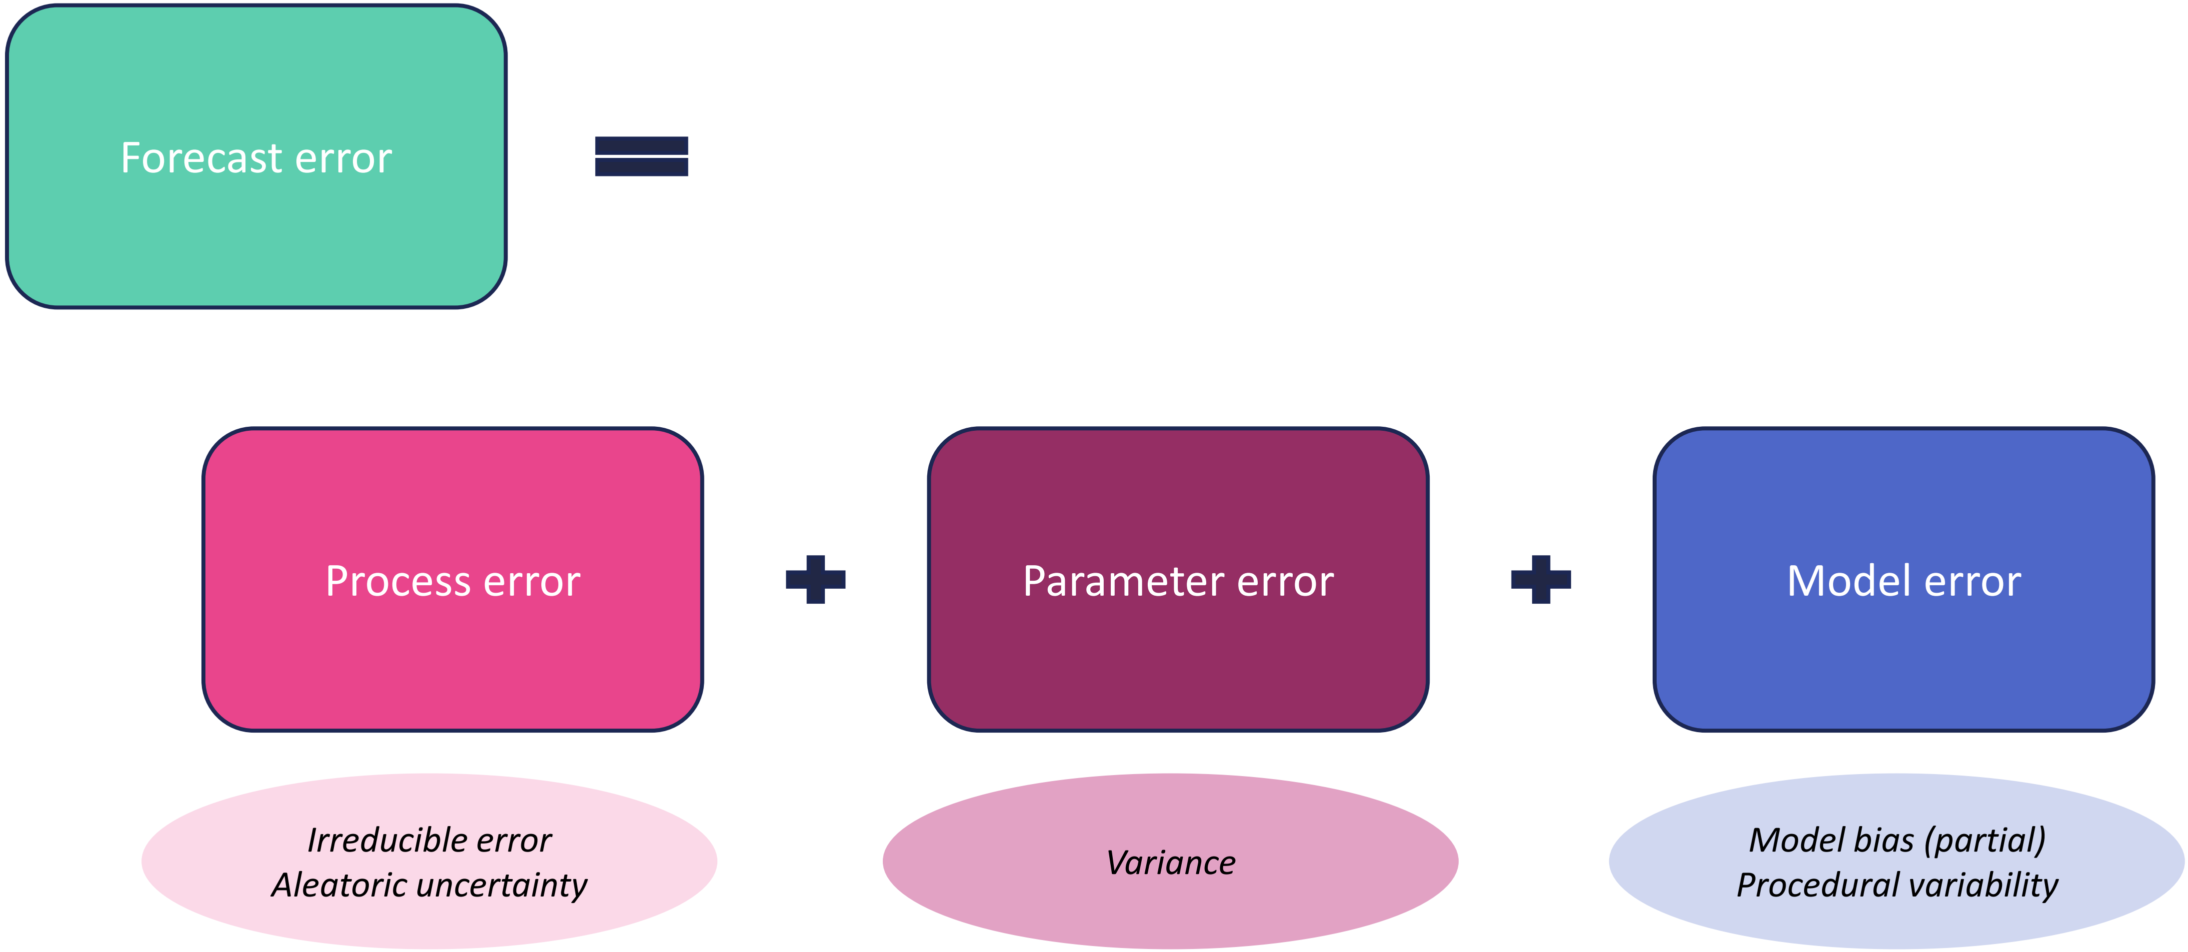

Divide and conquer some more

Forecast error = Process error + Parameter error + Model error

Model error = Model structure error + Model distribution error

Model error = Internal model structure error + External model structure error + Internal model distribution error + External model distribution error

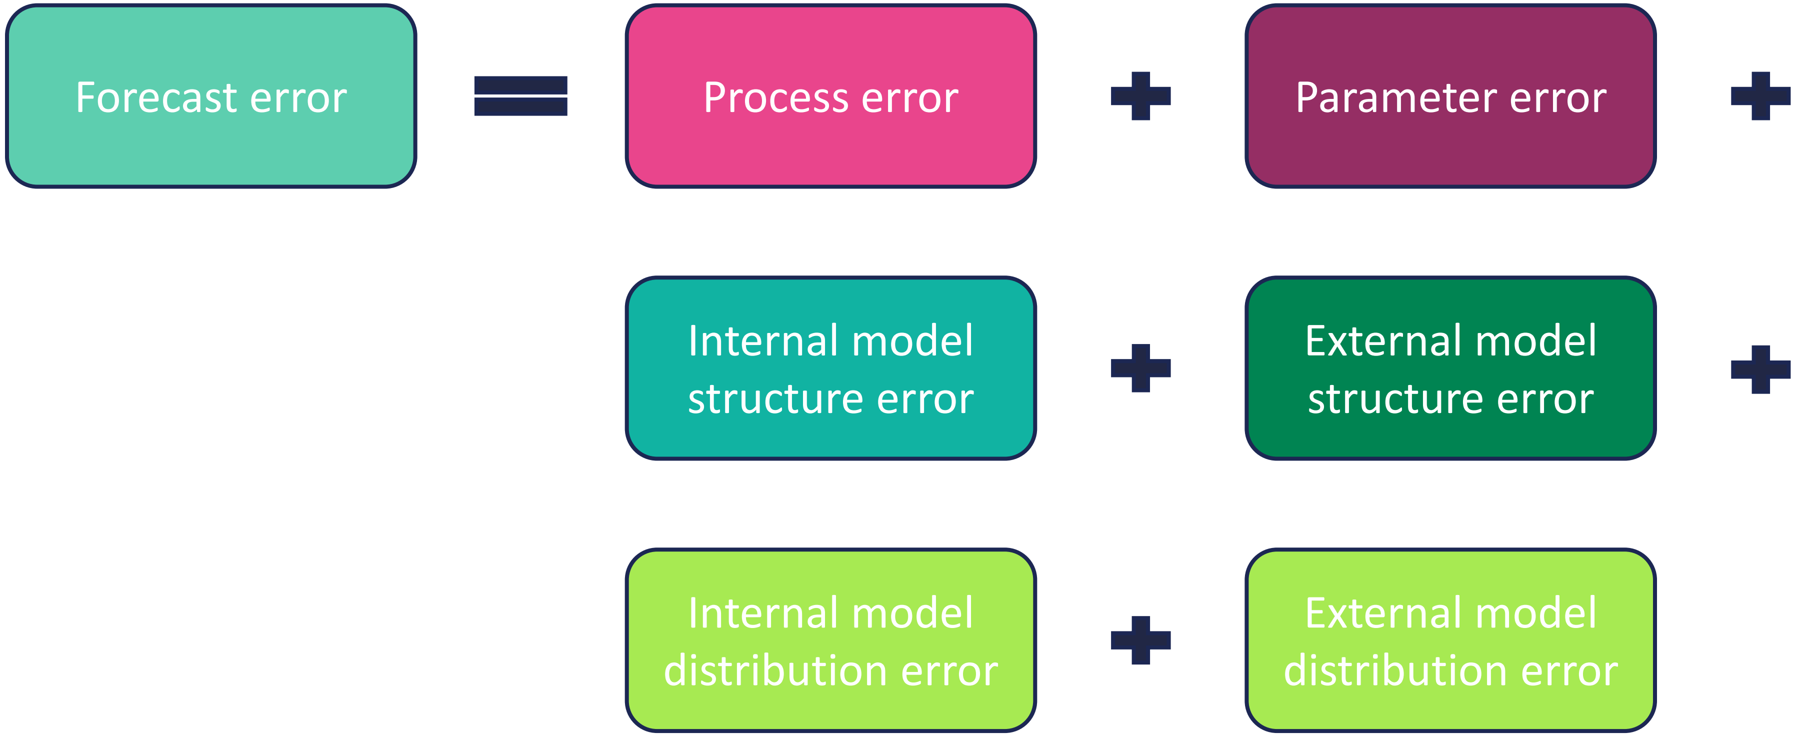
Forecast error
Process error
Parameter error
Internal model structure error
External model structure error
Internal model distribution error
External model distribution error

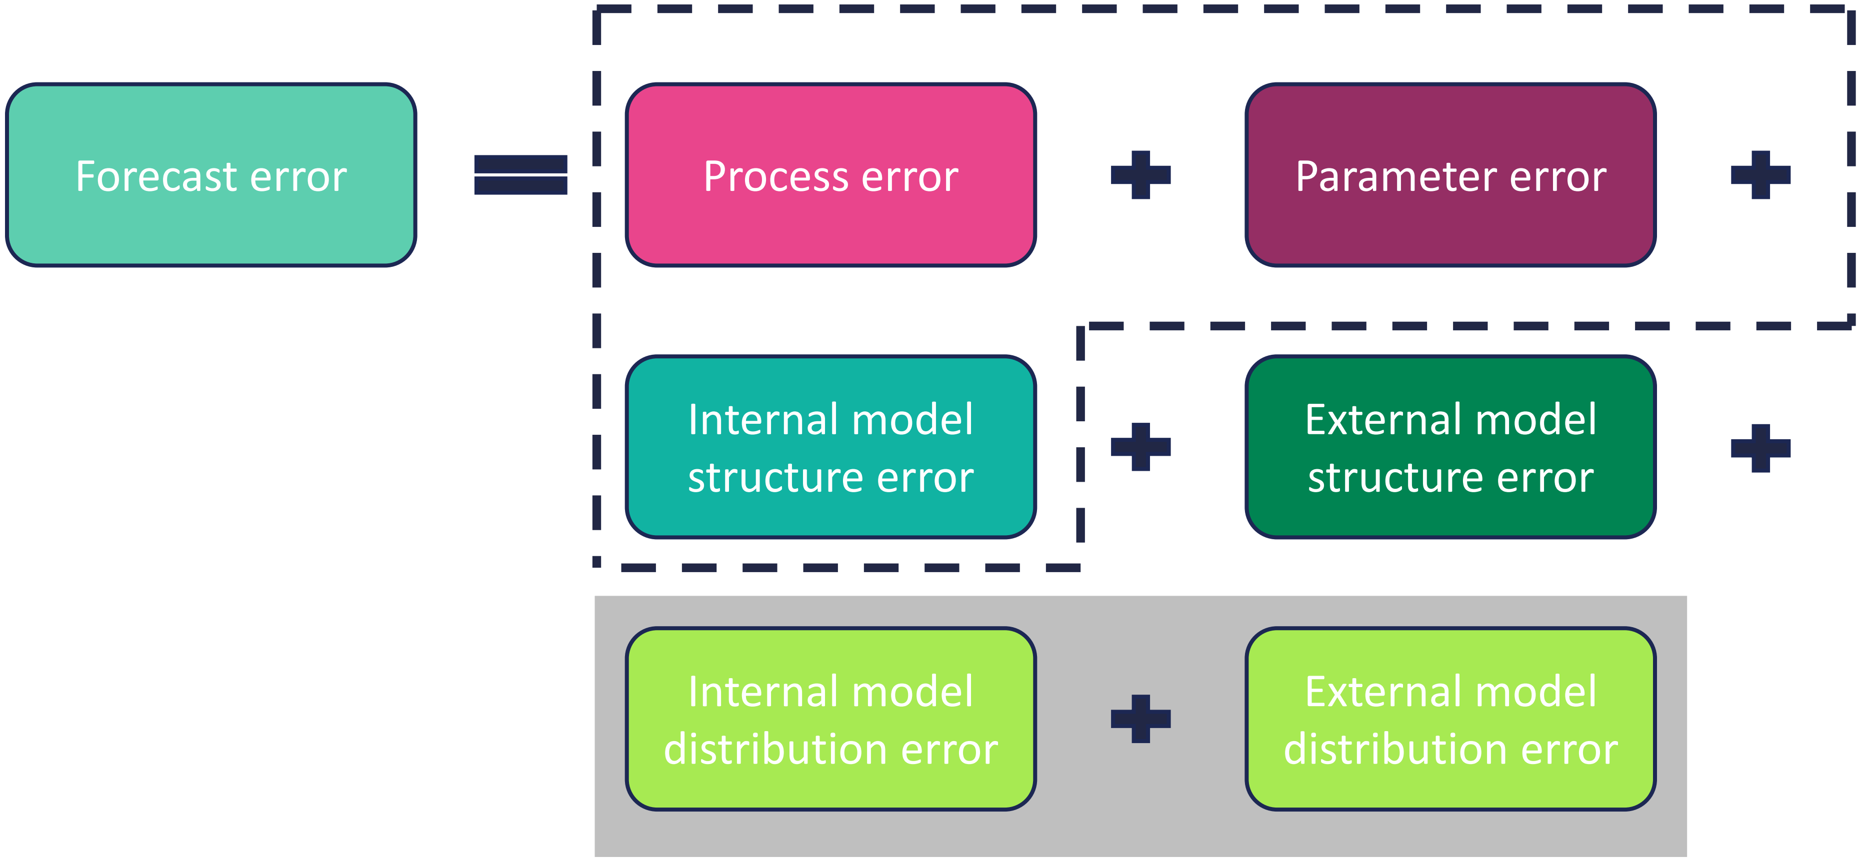
Forecast error
Process error
Parameter error
Internal model structure error
External model structure error
Internal model distribution error
External model distribution error

External model structure error

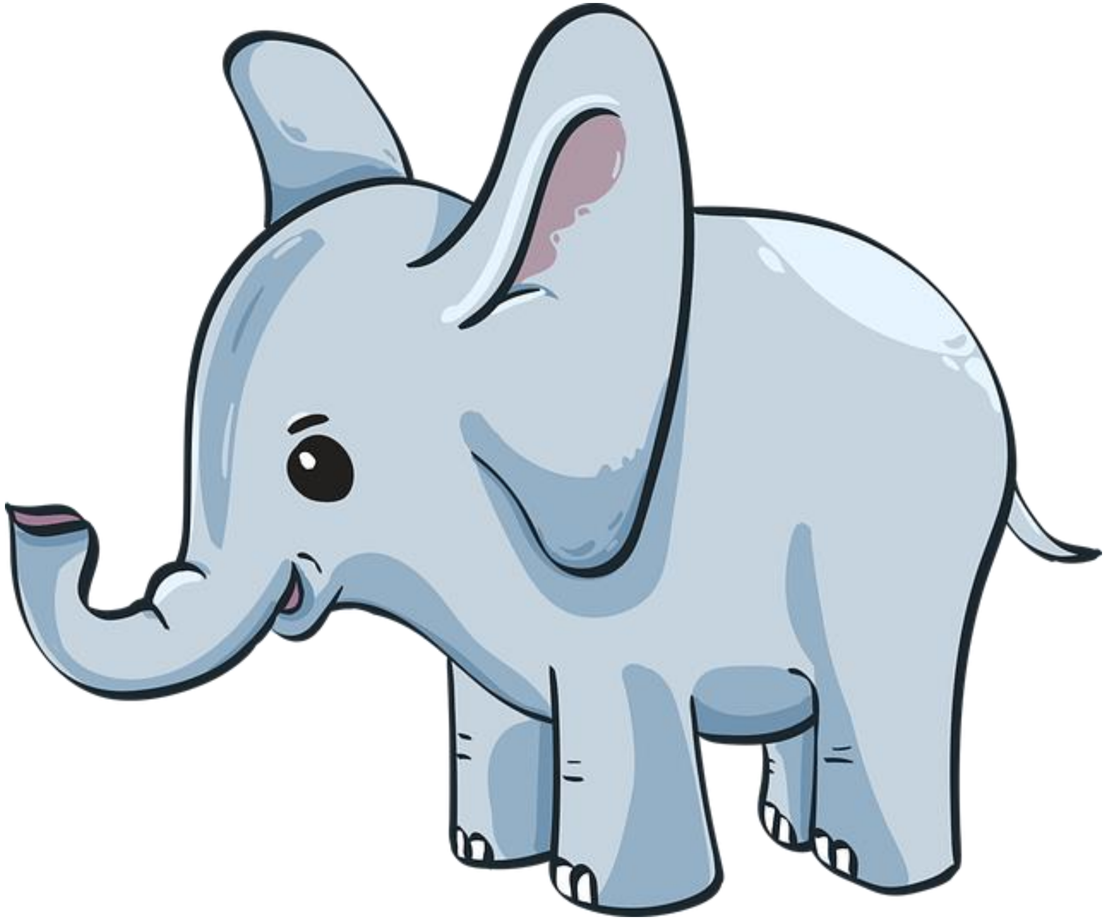

Estimating process error with Monte Carlo simulation

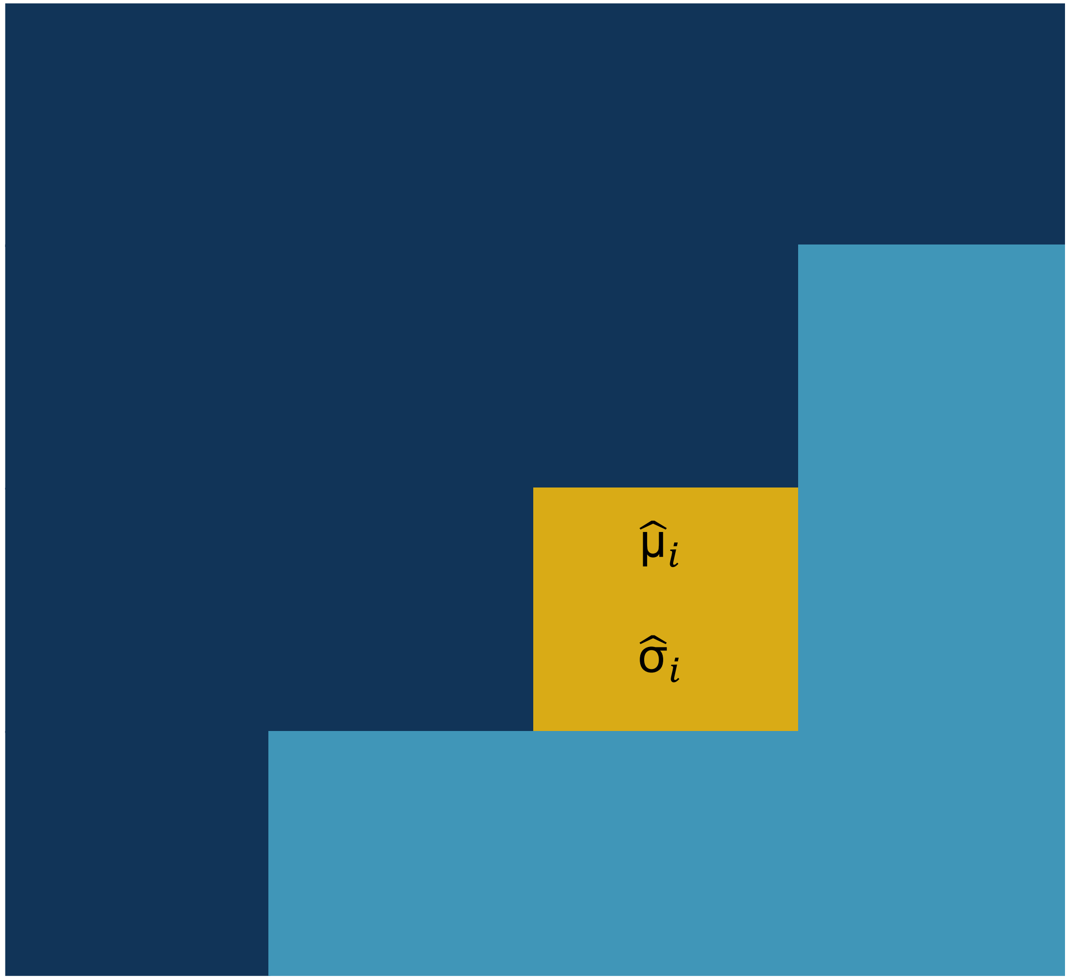

For each prediction, sample from

$$f(\hat{\mu}_i, \hat{\sigma}_i)$$

Parameter error with bootstrapping

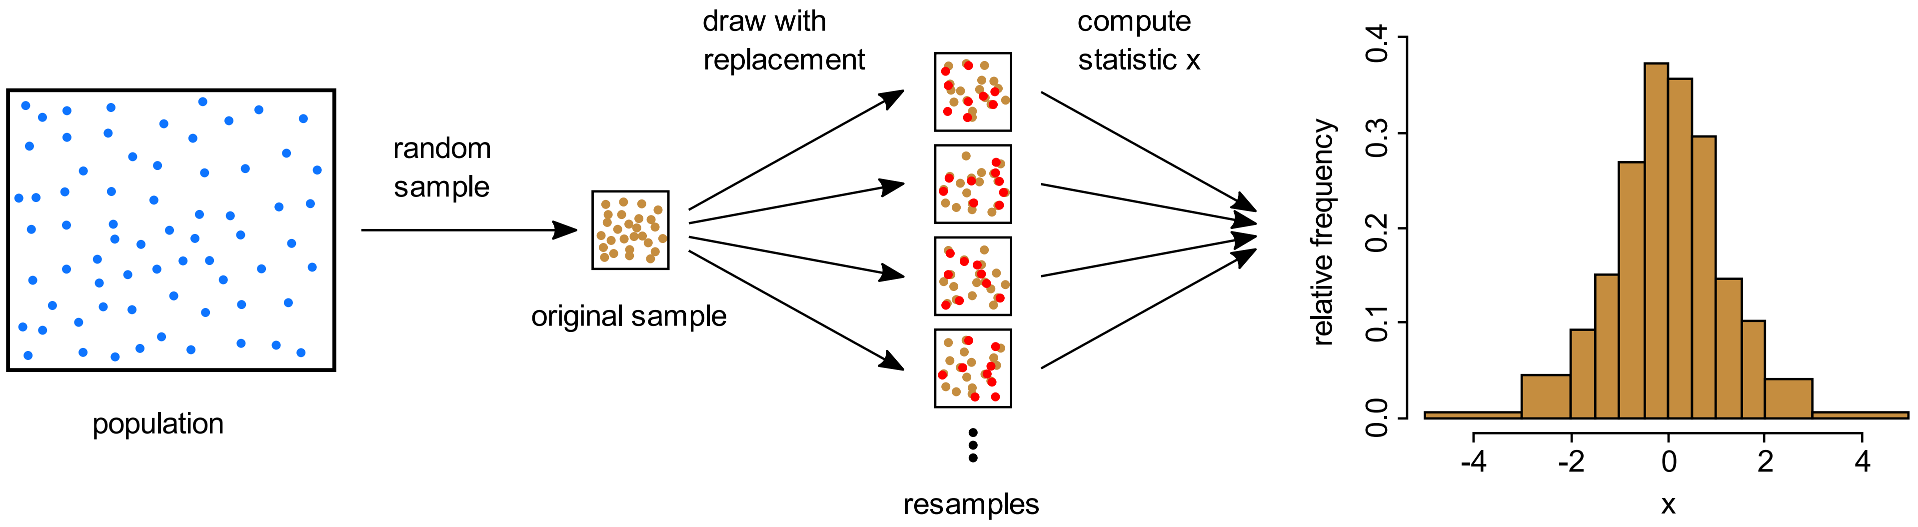

*By Biggerj1, Marsupilami - * File:Thist german.png, Autor: MM-Stathttps://postimg.cc/MffYNykZ, Autor Biggerj1, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=135426288*

Internal model structure error – Bayes and model sets

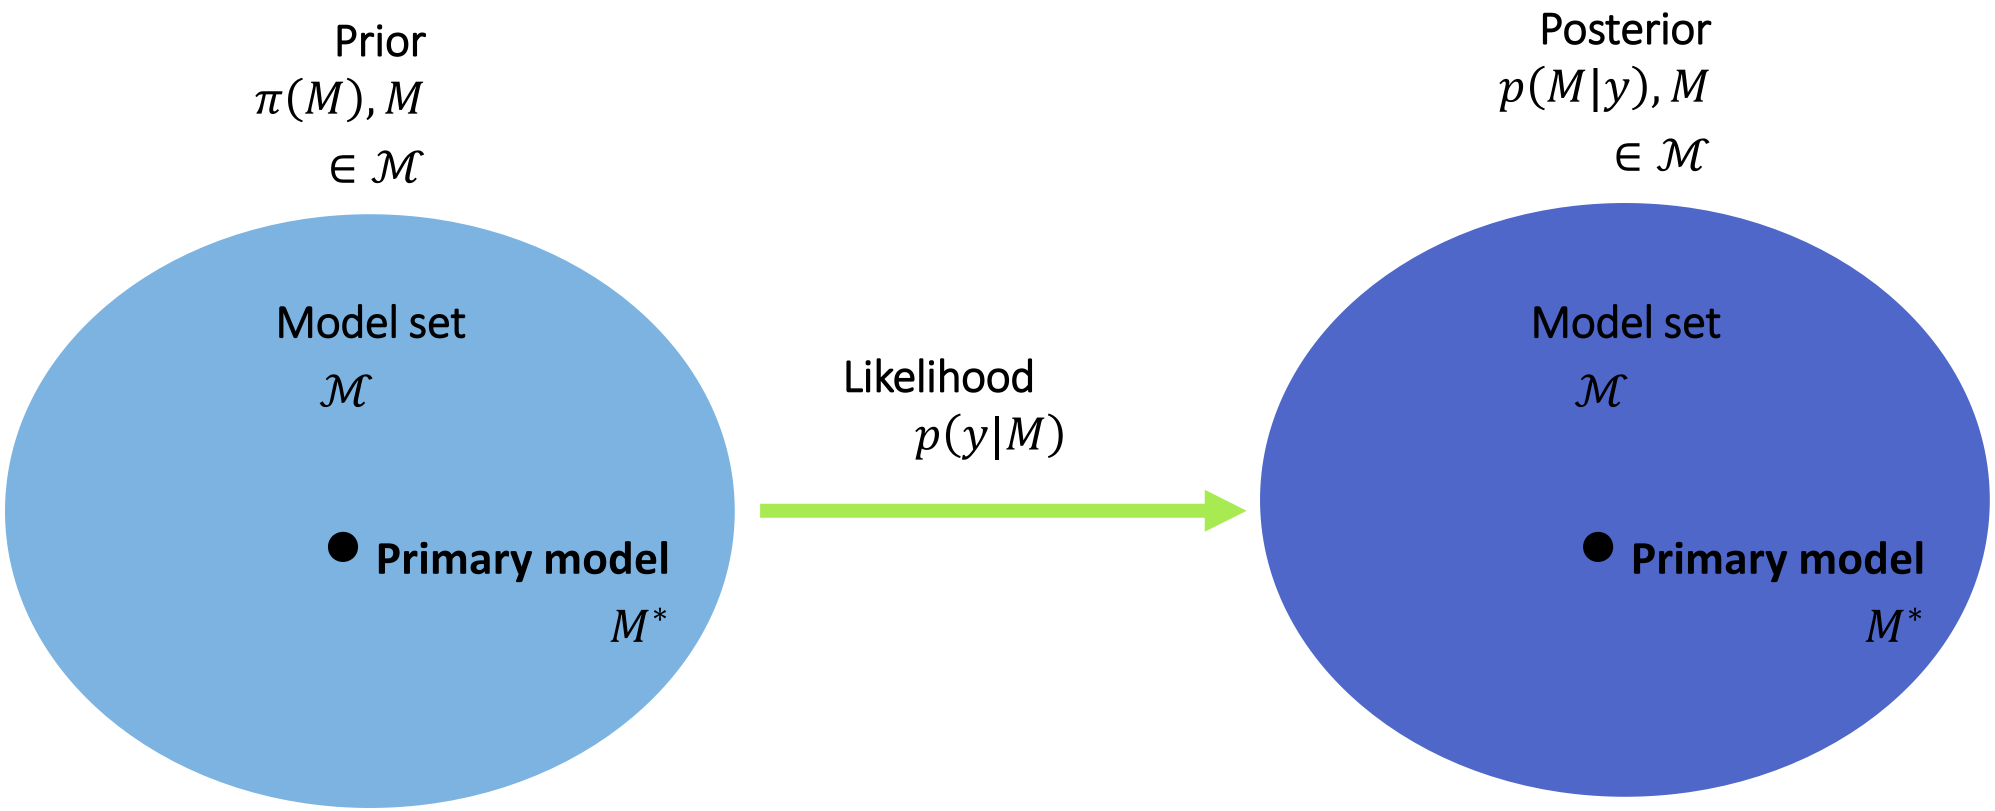

Secret recipe – remove some of the pie

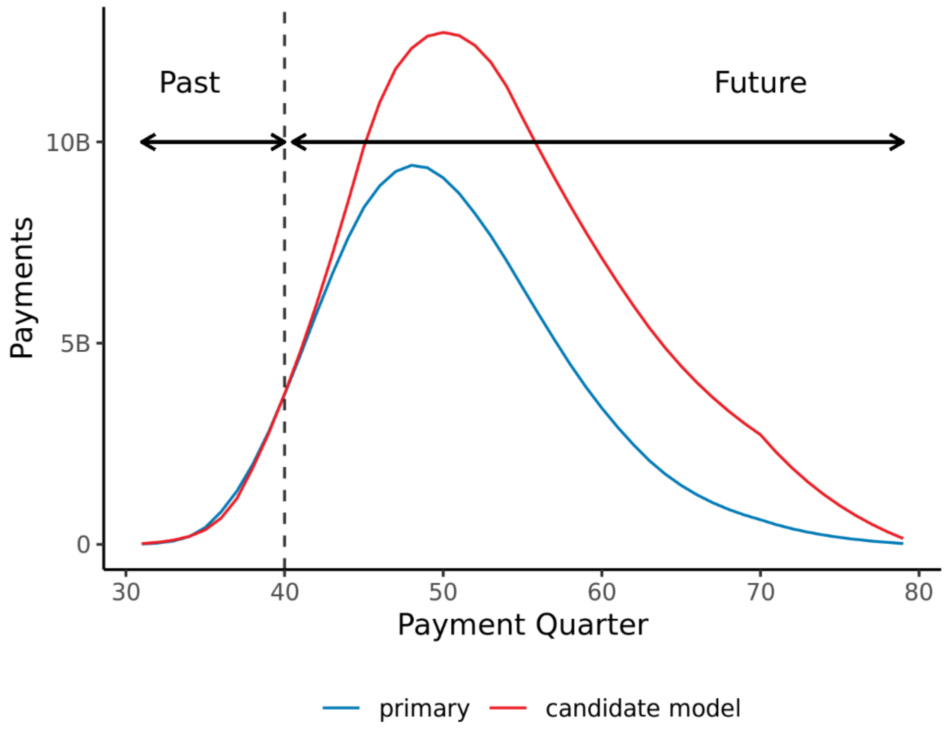

Some models match the past well
but go off the rails in the future

Posterior
$p(M|y), M \in \mathcal{M}$

Prune away models that don't extrapolate credibly

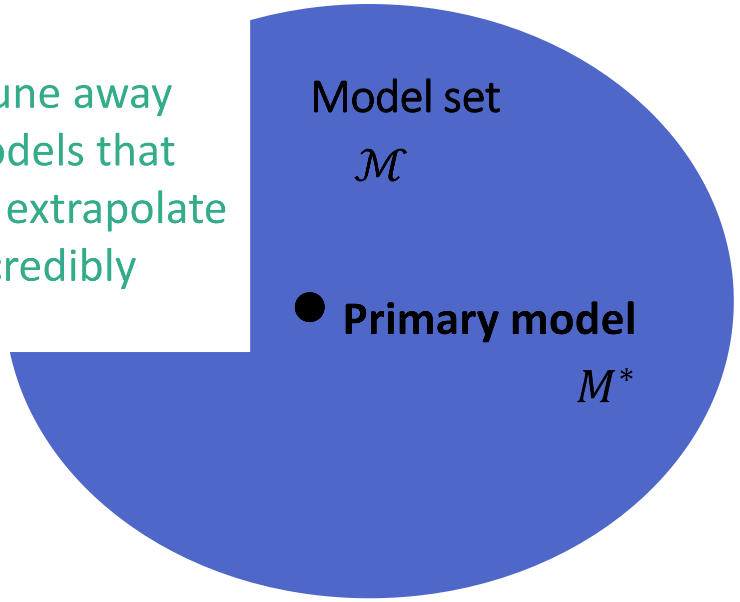

Lassoing the Model set

Bayesian lasso interpretation yields the model set

The Lasso

- Model form:
$$y = h^{-1}(X\beta) + \varepsilon$$
- Loss function
$$\widehat{\beta}(\lambda) = \arg\min_{\beta}[\ell(y|\beta) + \lambda^T|\beta|]$$
  - ℓ = negative log-likelihood **(NLL)**
  - $|.|$ operates elementwise on β
  - λ = **penalty parameter vector** with non-negative components
- Model set generation
  - Each λ corresponds to a different model
  - Path of models fitted for different λ

Prior distribution

- Laplace prior distribution
$$\pi(\beta) \propto exp(-\lambda^T|\beta|)$$

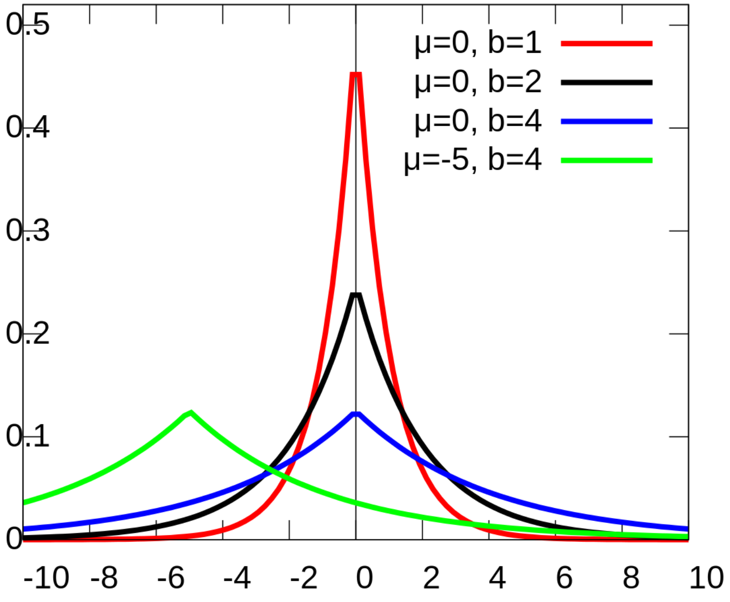

Specifying the prior

- Prior specification for a single parameter (excluding intercept):
  - Mean: 0
  - Variance: $Var[\beta_j] = 2/\lambda_j^2$
- All parameters (excluding intercept)
  - $\boldsymbol{\lambda}^{\boldsymbol{T}} = \lambda(1, \ldots, 1)$
  - λ = 0 → ML solution
  - λ → Inf →Intercept only model
  - What λ to use to lead to sensible model sets??

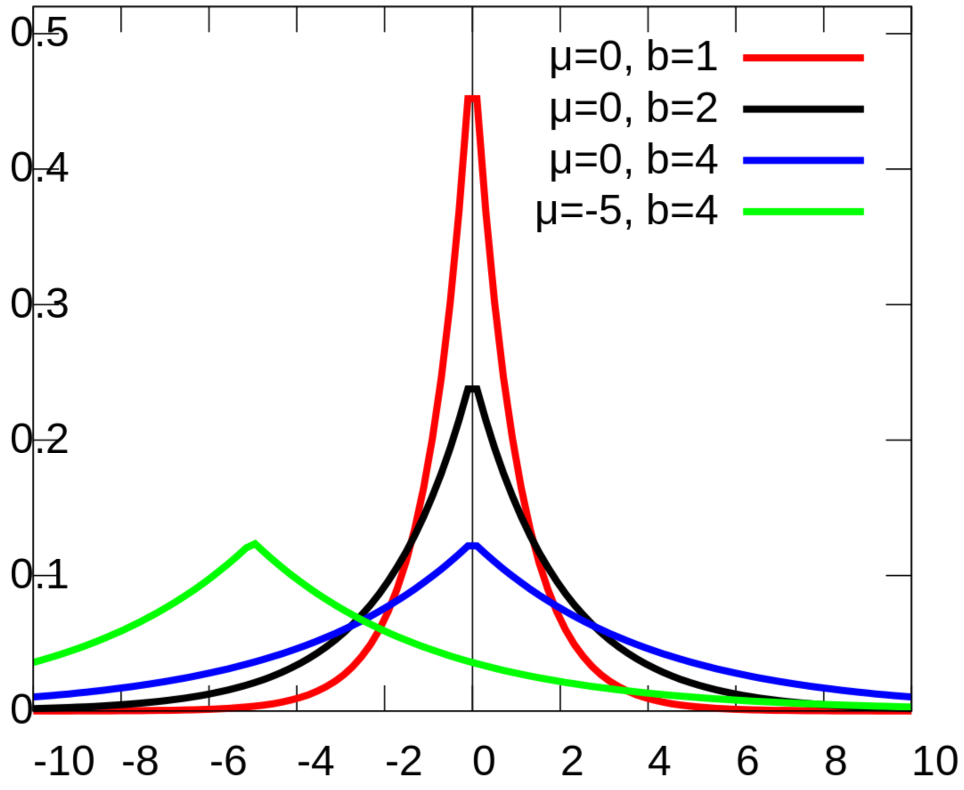

https://en.wikipedia.org/wiki/Laplace_distribution

Reasonable priors

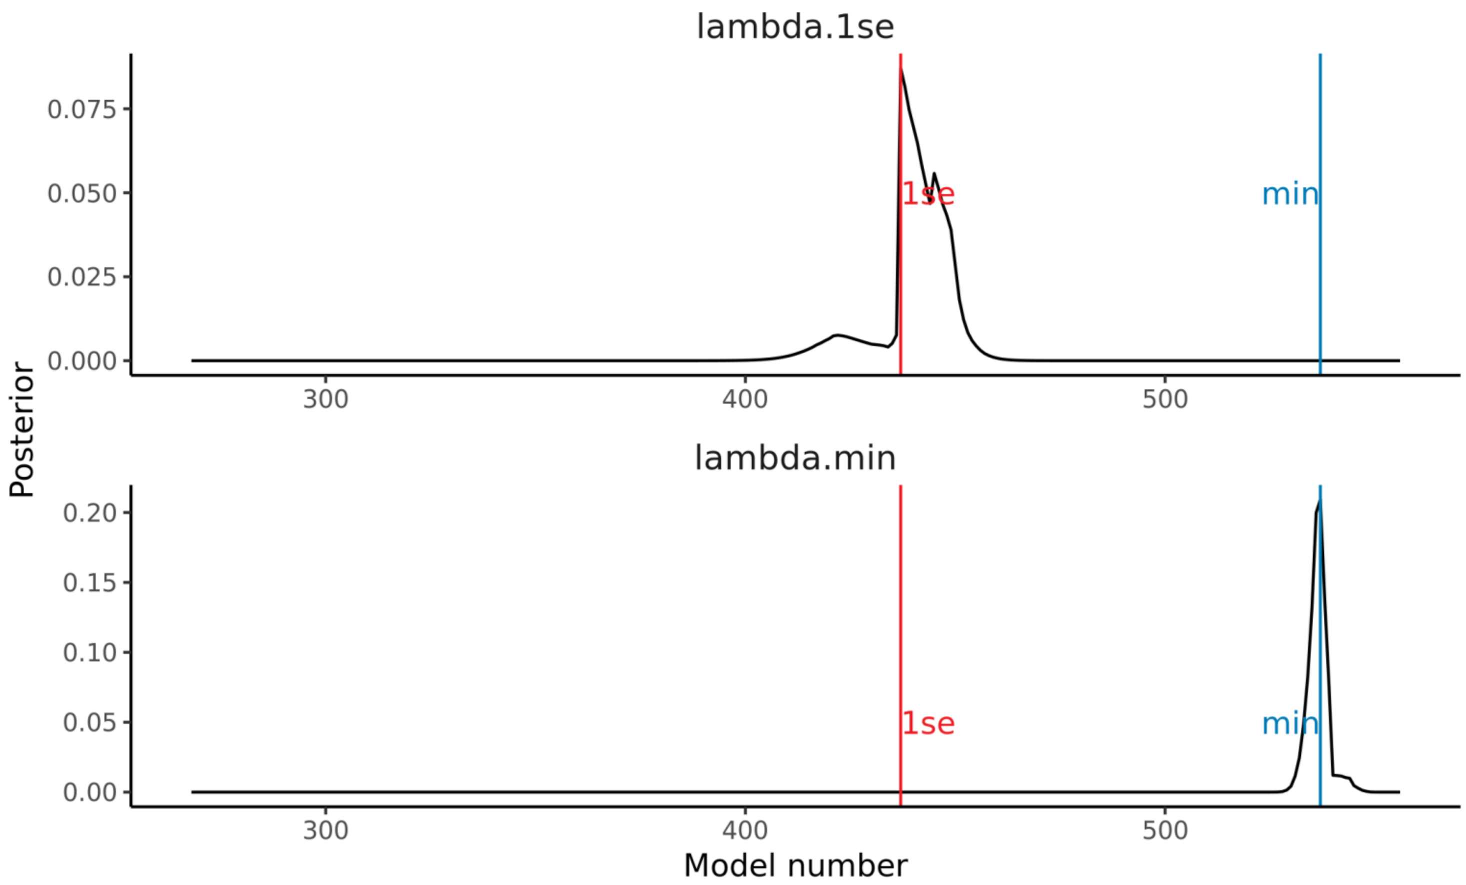

- Lasso models usually fitted using cross validation (CV) to select penalty to use
- Popular choices:
  - Lambda.min = penalty corresponding to model with minimum CV error
  - Lambda.1se = penalty where CV error 1 standard error from minimum – protects against over-fitting

Extreme (but reasonable?) priors

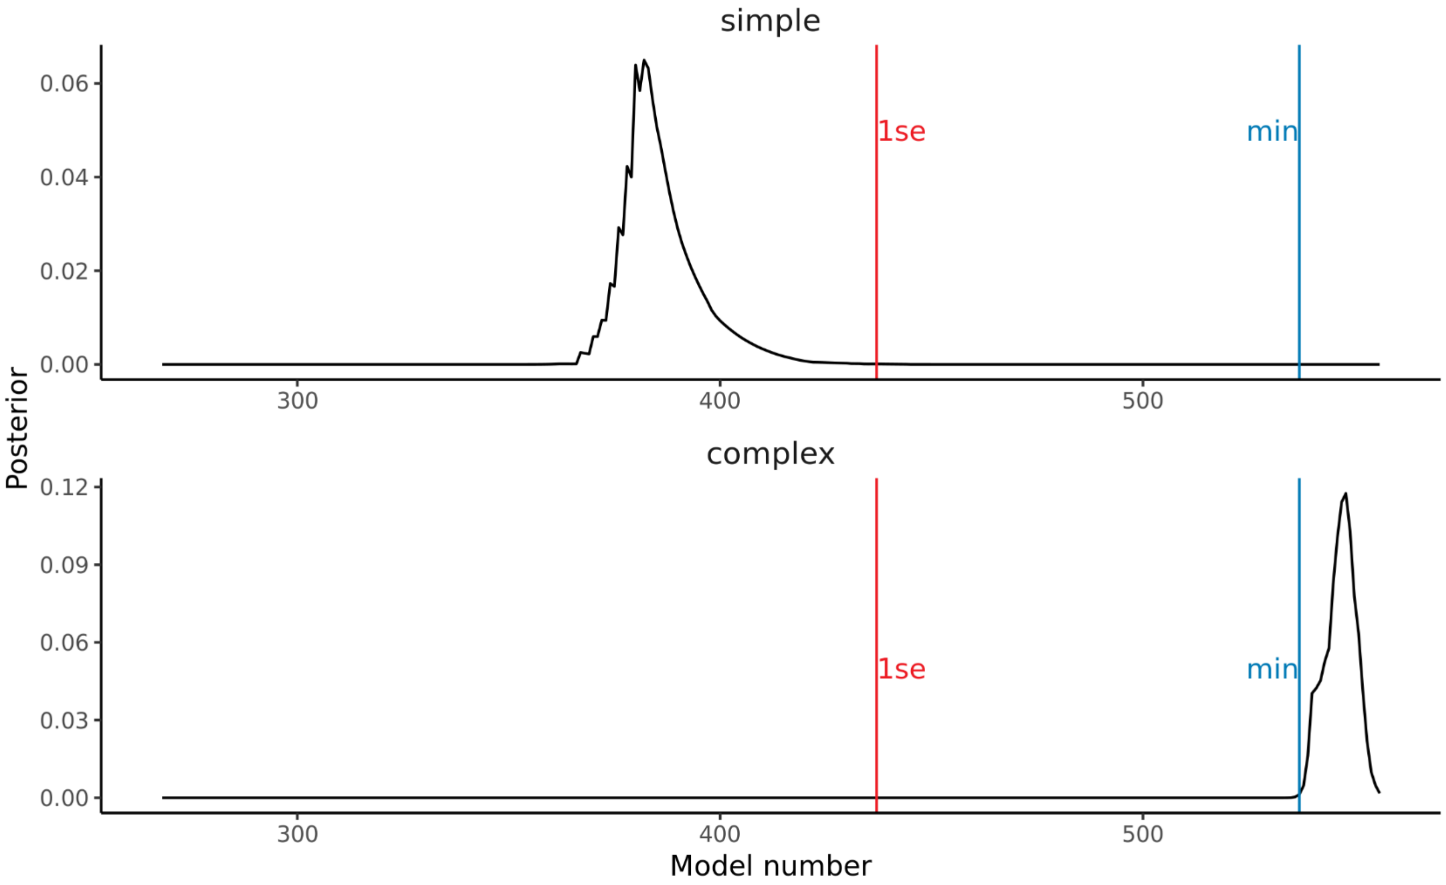

Synthetic data sets

- Data set 1
  - Satisfies chain ladder assumptions
- Data set 2
  - Payment period effect included
- Data set 3
  - Accident – development period interaction included for small number of recent cells
- Data set 4
  - Like data set 2 but payment period effect depends on development period

Internal model structure error

Data Set	LASSO Model	Loss Reserve			Estimated IMSE (CoV)
		True	Forecast		
			Raw 1se	Posterior	
		AUDB	AUDB	AUDB	%
1	Simple	190		198	0.7
	1se	190	194	194	0.4
	minCV	190		194	0.5
	Complex	190		203	0.8
2	Simple	238		260	0.1
	1se	238	261	260	0.1
	minCV	238		244	3.4
	Complex	238		272	3.1
3	Simple	608		877	1.7
	1se	608	778	777	6.8
	minCV	608		687	2.0
	Complex	608		875	5.8
4	Simple	216		244	0.2
	1se	216	247	247	0.3
	minCV	216		268	0.7
	Complex	216		276	1.2

- Model error estimated as variance over the model set
- Volatile estimates – “thin” posteriors
  - 10 – 30 models, not a lot
- Can we enhance with bootstrapping?
  - Also allows us to estimate parameter error + process error

Bootstrapping

Bootstrapping

Semi-parametric bootstrap

1. Primary model
2. Resample residuals
3. Pseudo data set
4. Refit model and estimate quantities of interest

Alternative: Non-parametric bootstrap

1. Stratified resampling of data
2. Pseudo data set
3. Refit model and estimate quantities of interest

Estimating uncertainty with the bootstrap

Pruned bootstrap!

Bootstrap	lambda				
	1	2	…	N	
1					mean, variance
2					mean, variance
.					
.					
B					mean, variance

Separate Monte Carlo simulation for process error

Internal model error:
Mean of the row variances

Parameter error:
Variance of the row means

Results

Numerical results

Data set	Prior	Forecast					
		True ($B)	Mean ($B)	Internal model error (CoV)	Parameter error (CoV)	Process error (CoV)	Total error (CoV)
1	1se	190	189	0.32%	5.30%	3.29%	6.24%
	lambda.min	190	192	0.41%	5.15%	2.75%	5.85%
2	1se	238	252	1.45%	10.00%	3.93%	10.84%
	lambda.min	238	240	1.79%	8.83%	4.69%	10.16%
3	1se	608	703	2.27%	11.23%	5.71%	12.80%
	lambda.min	608	589	2.12%	11.19%	5.27%	12.54%
4	1se	216	243	1.37%	8.63%	4.01%	9.62%
	lambda.min	216	252	1.81%	12.54%	5.08%	13.65%

Conclusion

Take-homes

Components of variability

Machine learning and potential to measure model error

Technical approach to reserving including variability estimation

Pragmatic bootstrapping tips

Other comments

External model error

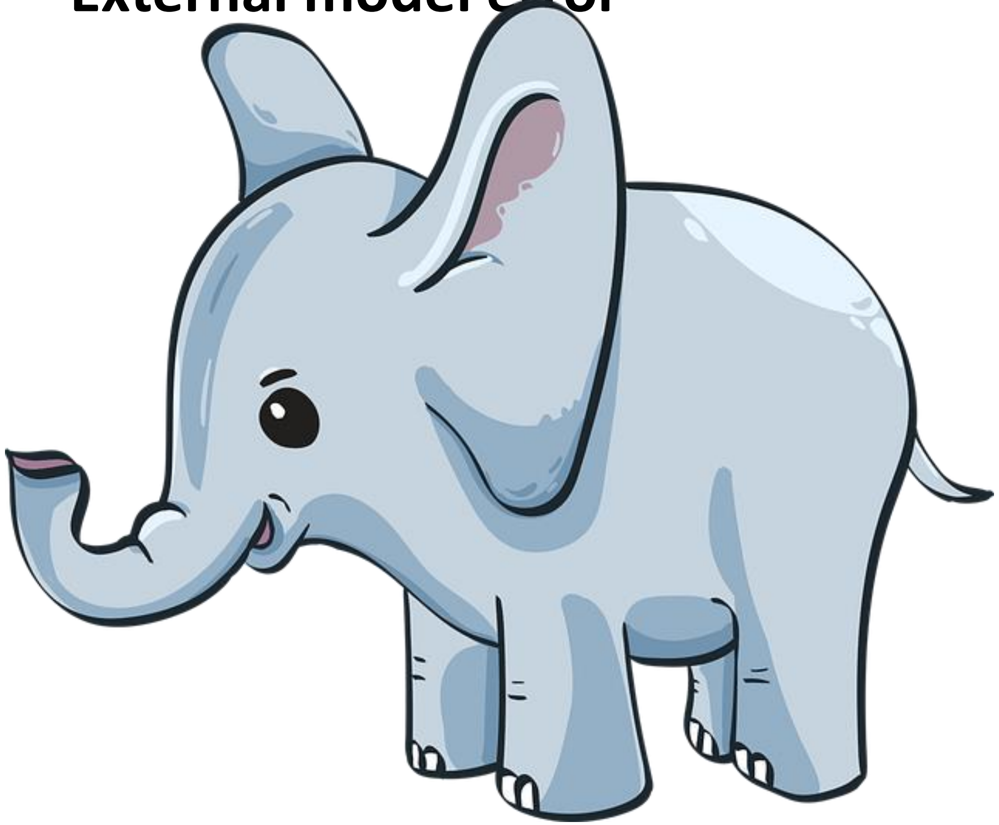

Model and parameter error are linked

Internal model error leaks into parameter error – so consider combined estimate only

Questions

Comments

The views expressed in this presentation are those of the presenter.