

# Approximate Bayesian Computation in Insurance

Patrick J. Laub and Pierre-Olivier Goffard



LABORATOIRE  
**SAF**  
SCIENCES ACTUARIELLE  
& FINANCIÈRE

# Motivation

Have a random number of claims  $N \sim p_N(\cdot; \boldsymbol{\theta}_{\text{freq}})$  and the claim sizes  $U_1, \dots, U_N \sim f_U(\cdot; \boldsymbol{\theta}_{\text{sev}})$ .

We aggregate them somehow, like:

- aggregate claims:  $X = \sum_{i=1}^N U_i$
- maximum claims:  $X = \max_{i=1}^N U_i$
- stop-loss:  $X = (\sum_{i=1}^N U_i - c)_+$ .

**Question:** Given a sample  $X_1, \dots, X_n$  of the summaries, what is the  $\boldsymbol{\theta} = (\boldsymbol{\theta}_{\text{freq}}, \boldsymbol{\theta}_{\text{sev}})$  which explains them?

E.g. a reinsurance contract

# Likelihoods

For simple rv's we know their likelihood (normal, exponential, gamma, etc.).

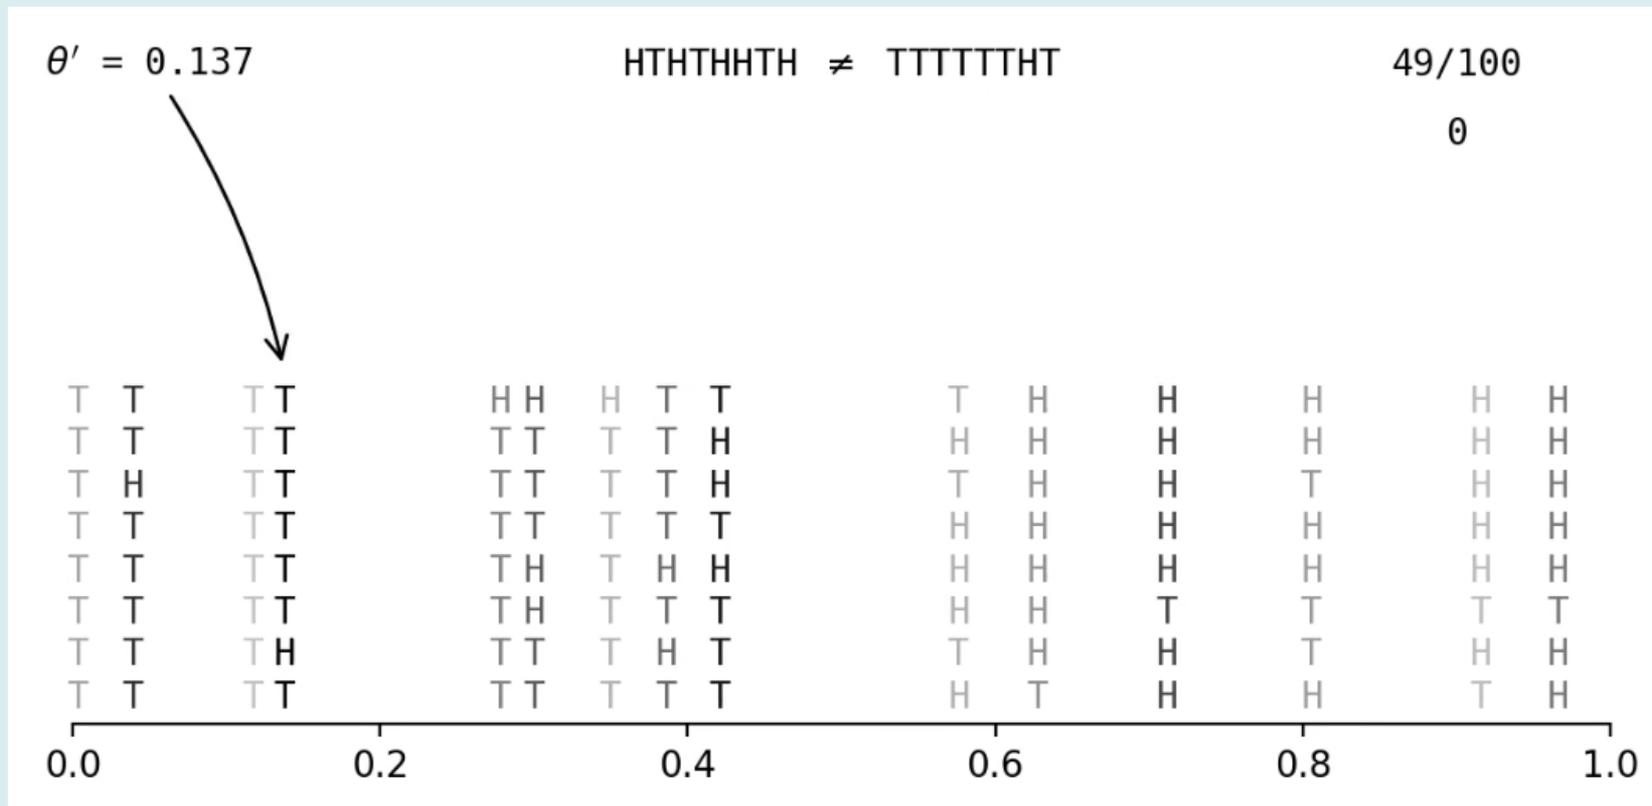
When simple rv's are combined, the resulting thing rarely has a tractable likelihood.

$$X_1, X_2 \stackrel{\text{i.i.d.}}{\sim} f_X(\cdot) \Rightarrow X_1 + X_2 \sim \text{Intractable likelihood!}$$

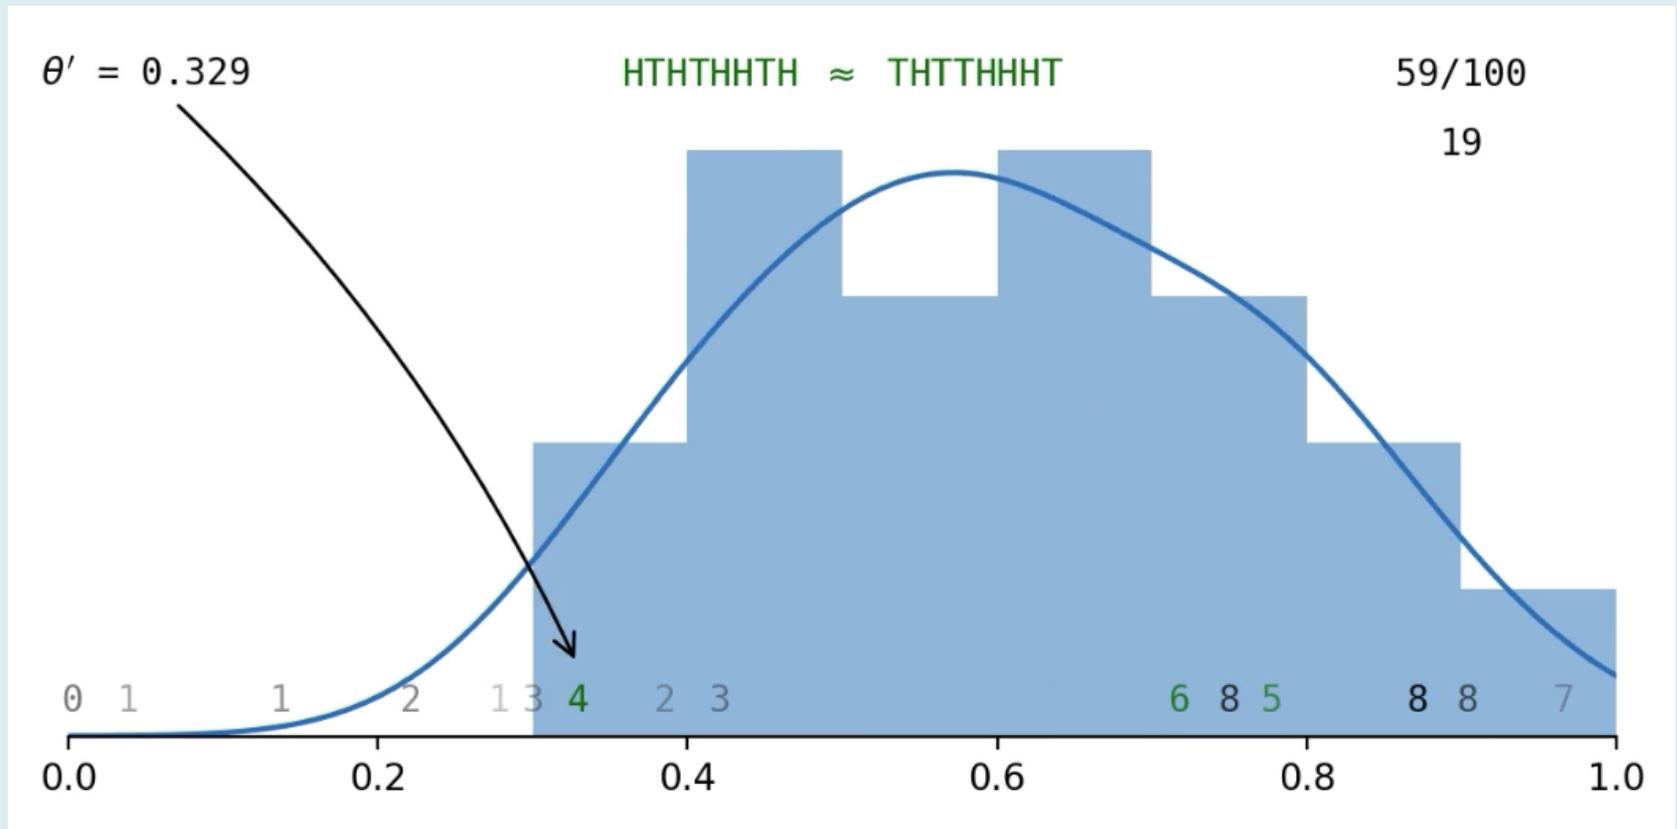
*Usually it's still possible to simulate these things...*



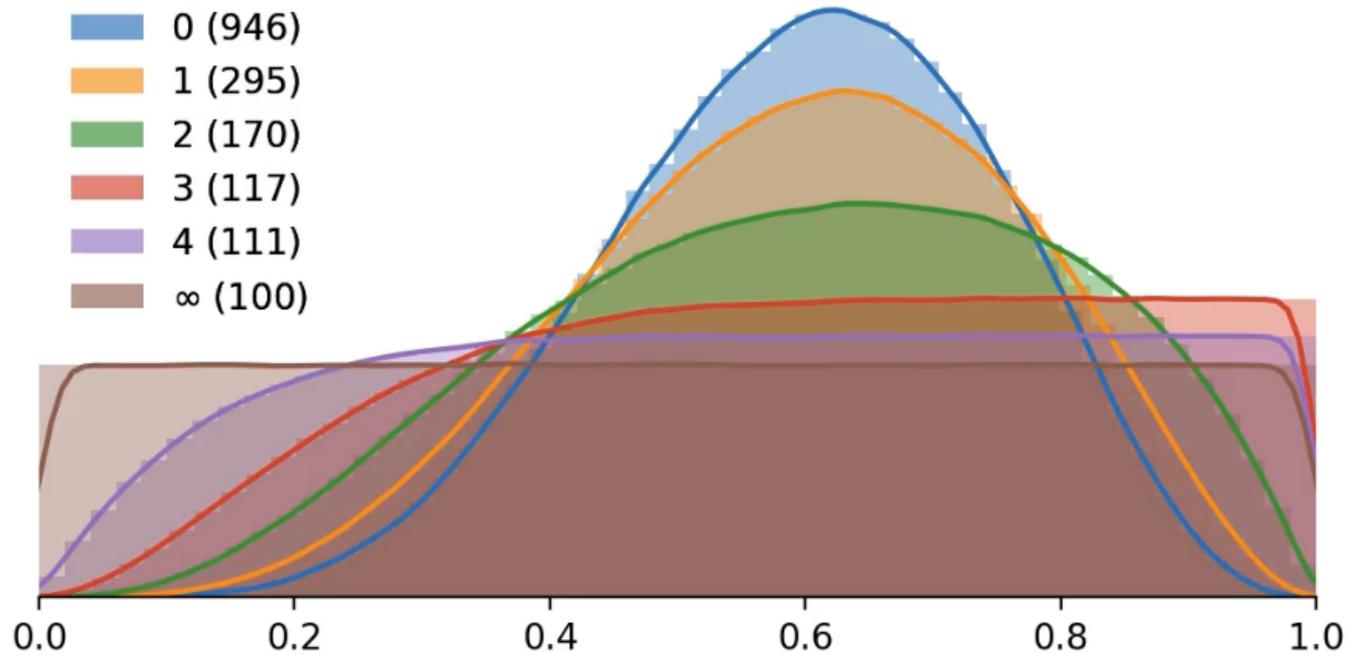
# Getting an exact match of the data is hard...



# Accept fake data that's close to observed data



# The 'approximate' part of ABC



# Does it work in theory?

We sample the *approximate / ABC posterior*

$$\pi_\epsilon(\boldsymbol{\theta} \mid \mathbf{x}_{\text{obs}}) \propto \pi(\boldsymbol{\theta}) \times \mathbb{P}(\|\mathbf{x}_{\text{obs}} - \mathbf{x}^*\| \leq \epsilon \text{ where } \mathbf{x}^* \sim \boldsymbol{\theta}),$$

but we care about the true posterior  $\pi(\boldsymbol{\theta} \mid \mathbf{x}_{\text{obs}})$ .

**Proposition:** Say we have continuous data  $\mathbf{x}_{\text{obs}}$ , and our prior  $\pi(\boldsymbol{\theta})$  has bounded support.

If for some  $\epsilon > 0$  we have

$$\sup_{(\mathbf{z}, \boldsymbol{\theta}) : \mathcal{D}(\mathbf{z}, \mathbf{x}_{\text{obs}}) < \epsilon, \boldsymbol{\theta} \in \Theta} \pi(\mathbf{z} \mid \boldsymbol{\theta}) < \infty$$

then for each  $\boldsymbol{\theta} \in \Theta$

$$\lim_{\epsilon \rightarrow 0} \pi_\epsilon(\boldsymbol{\theta} \mid \mathbf{x}_{\text{obs}}) = \pi(\boldsymbol{\theta} \mid \mathbf{x}_{\text{obs}}).$$

# Mixed data

We filled in some of the blanks for *mixed data*.

Our data was mostly continuous data but had an atom at 0.

Get

$$\lim_{\epsilon \rightarrow 0} \pi_{\epsilon}(\boldsymbol{\theta} \mid \mathbf{x}_{\text{obs}}) = \pi(\boldsymbol{\theta} \mid \mathbf{x}_{\text{obs}})$$

when

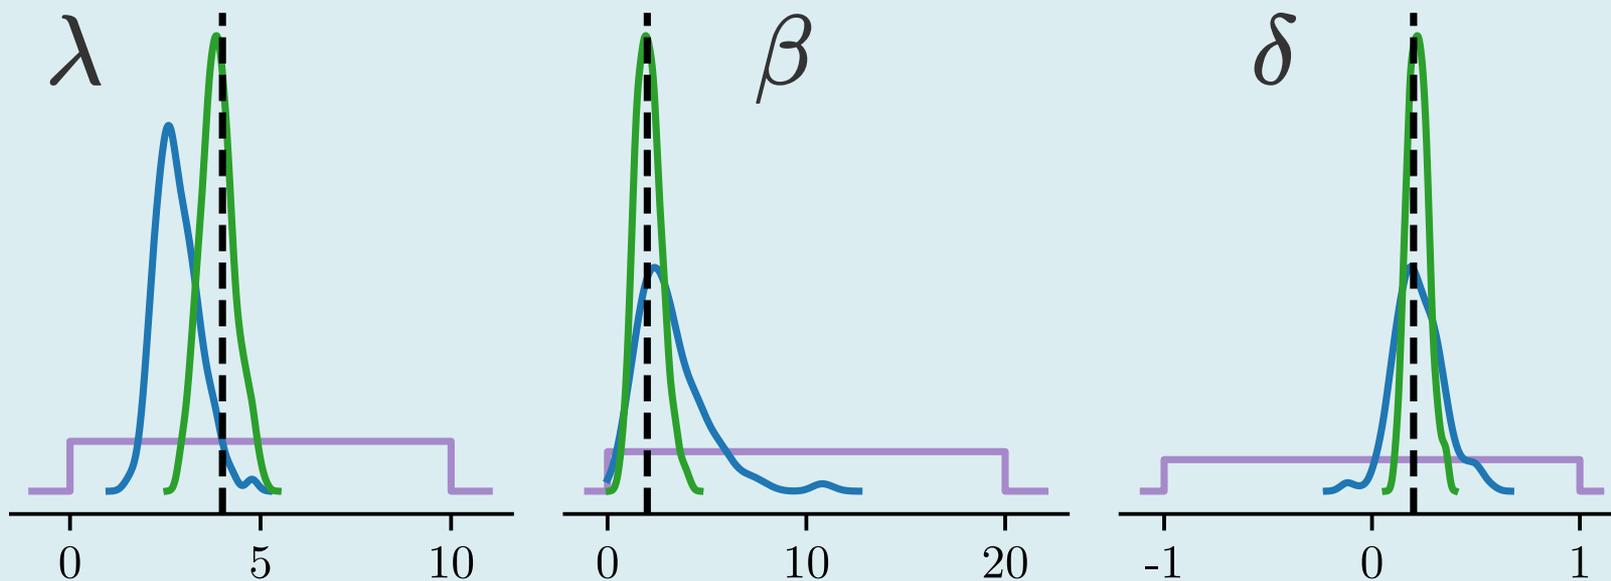
$$\mathcal{D}(\mathbf{z}, \mathbf{x}_{\text{obs}}) = \begin{cases} \mathcal{D}^+(\mathbf{z}^+, \mathbf{x}_{\text{obs}}^+) & \text{if } \#\text{Zeros}(\mathbf{z}) = \#\text{Zeros}(\mathbf{x}_{\text{obs}}), \\ \infty & \text{otherwise.} \end{cases}$$

# Does it work in practice?

In other words, when does it break & how slow is it?

# Dependent example

- frequency  $N \sim \text{Poisson}(\lambda = 4)$ ,
- severity  $U_i \mid N \sim \text{DepExp}(\beta = 2, \delta = 0.2)$ , defined as  $U_i \mid N \sim \text{Exp}(\beta \times e^{\delta N})$ ,
- summation summary  $X = \sum_{i=1}^N U_i$



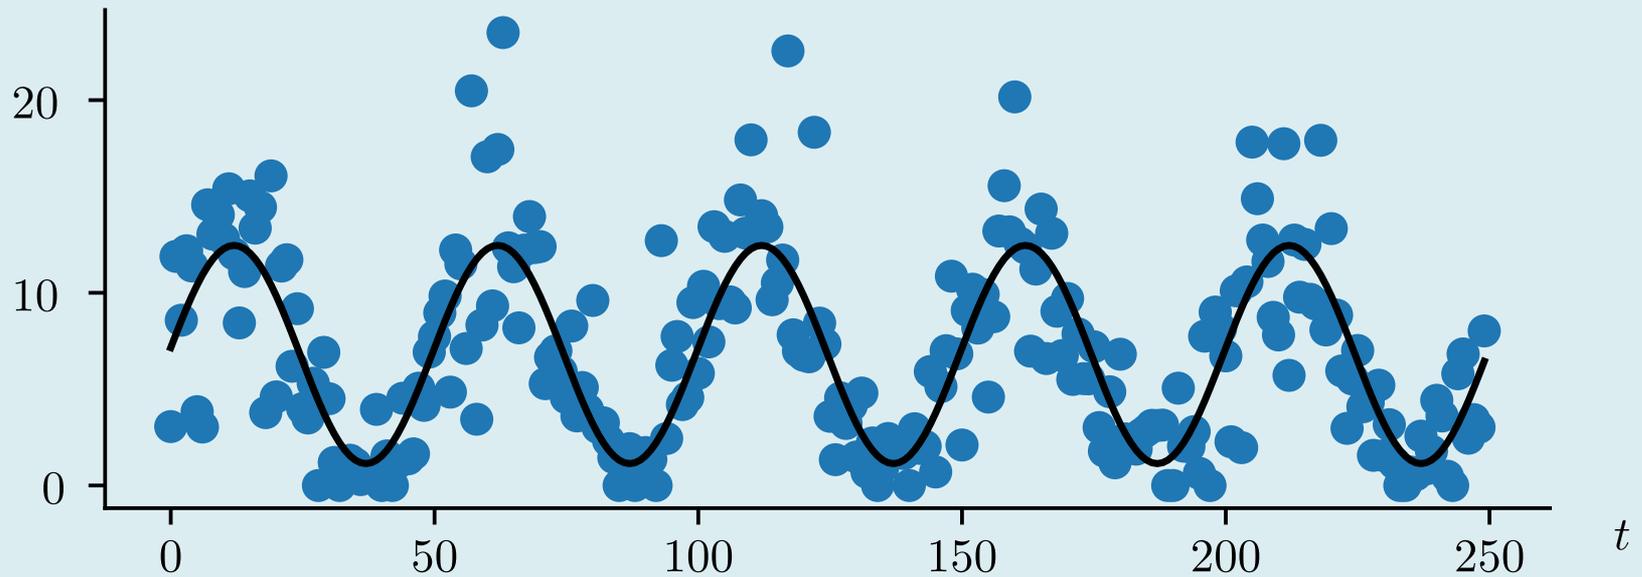
ABC posteriors based on 50  $X$ 's and on 250  $X$ 's given uniform priors.

# pip install approxbayescomp

```
1 import approxbayescomp as abc
2
3 # Load data to fit
4 obsData = ...
5
6 # Frequency-Loss Model
7 freq = "poisson"
8 sev = "frequency dependent exponential"
9 psi = abc.Psi("sum") # Aggregation process
10
11 # Fit the model to the data using ABC
12 prior = abc.IndependentUniformPrior([(0, 10), (0, 20), (-1, 1)])
13 model = abc.Model(freq, sev, psi, prior)
14 fit = abc.smc(numIters, popSize, obsData, model)
```

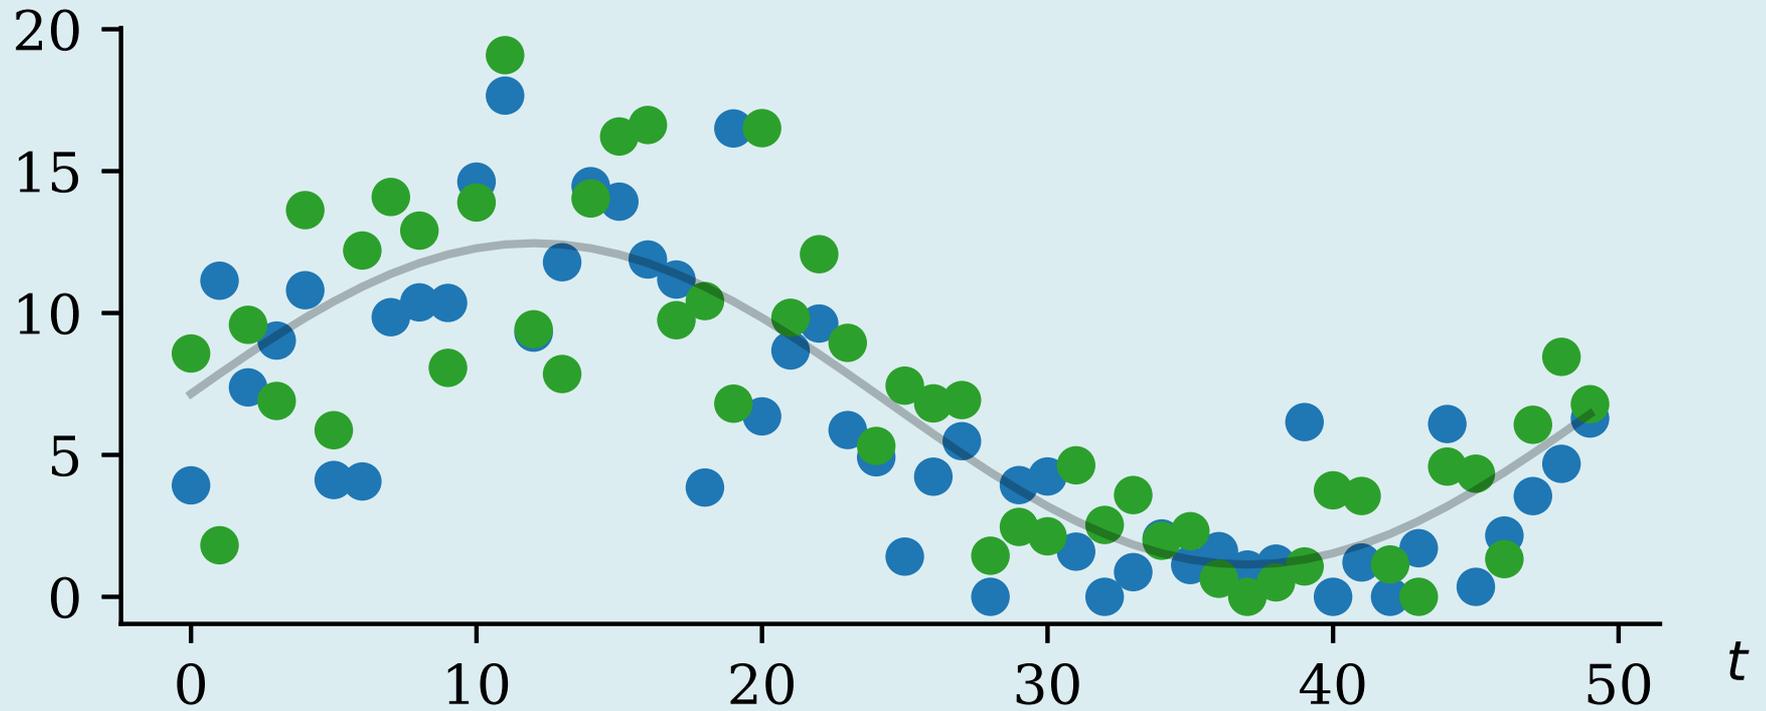
# Time-varying example

- Claims form a Poisson process point with arrival rate  $\lambda(t) = a + b[1 + \sin(2\pi ct)]$ .
- The observations are  $X_s = \sum_{i=N_{s-1}}^{N_s} U_i$ .
- Frequencies are  $\text{CPoisson}(a = 1, b = 5, c = \frac{1}{50})$  and sizes are  $U_i \sim \text{Lognormal}(\mu = 0, \sigma = 0.5)$ .



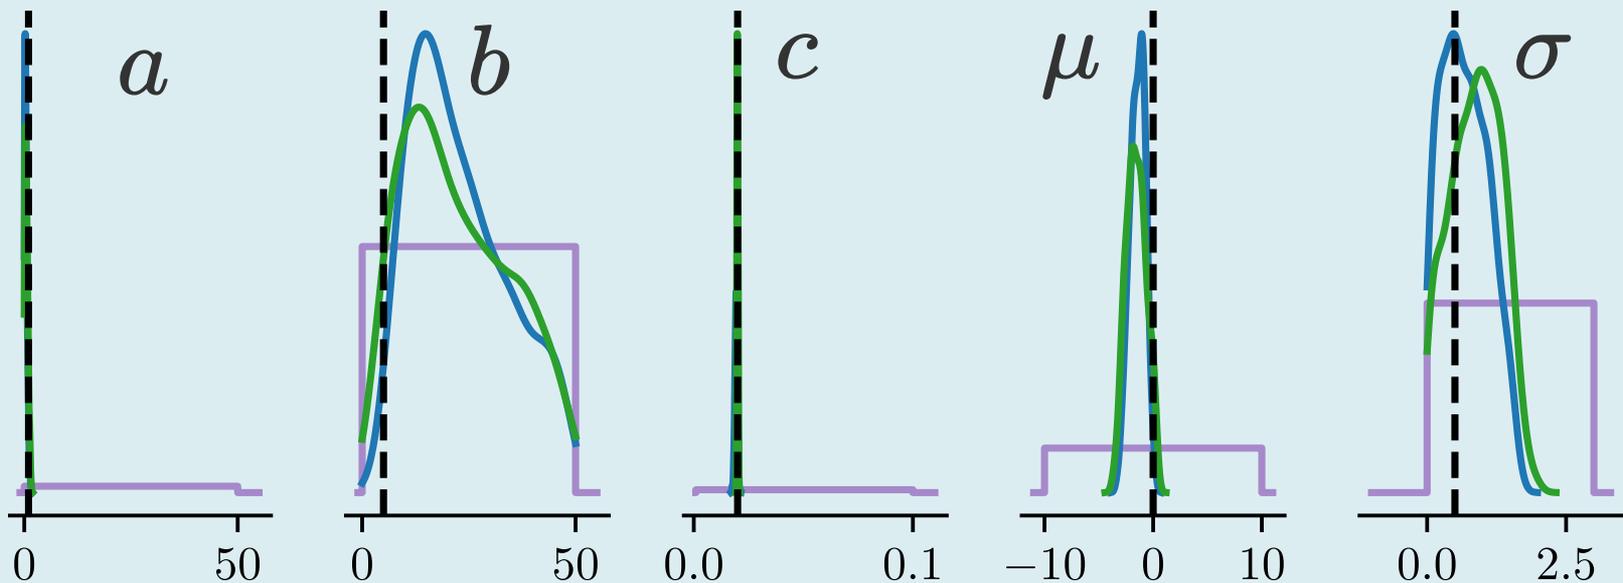
— Expected value for  $x_s$       ● Observed  $x_s$

From the same  $\theta$  the data is **random**



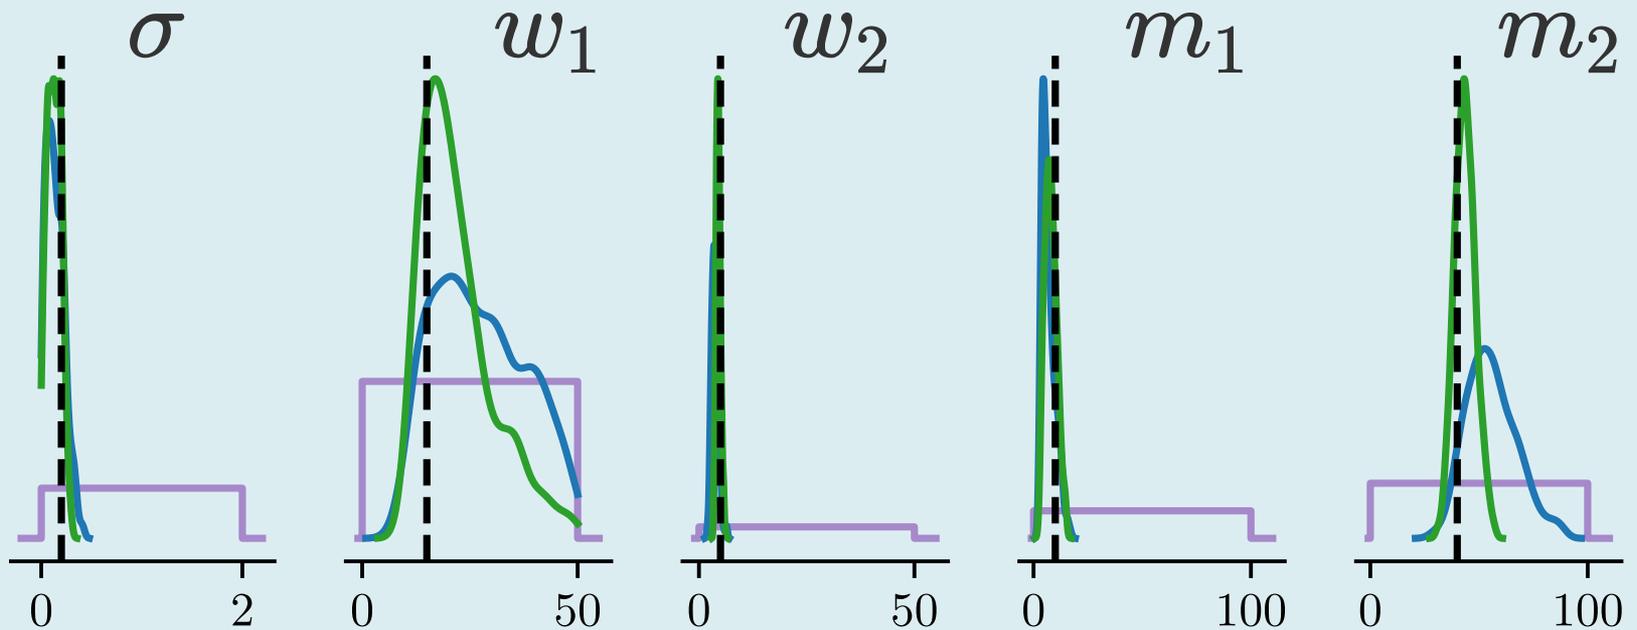
# Time-varying example

- Claims form a Poisson process point with arrival rate  $\lambda(t) = a + b[1 + \sin(2\pi ct)]$ .
- The observations are  $X_s = \sum_{i=N_{s-1}}^{N_s} U_i$ .
- Frequencies are  $\text{CPoisson}(a = 1, b = 5, c = \frac{1}{50})$  and sizes are  $U_i \sim \text{Lognormal}(\mu = 0, \sigma = 0.5)$ .
- ABC posteriors based on 50  $X$ 's and on 250  $X$ 's with uniform priors.



# Bivariate example

- Two lines of business with dependence in the claim frequencies.
- Say  $\Lambda_i \sim \text{Lognormal}(\mu \equiv 0, \sigma = 0.2)$ .
- Claim frequencies  $N_i \sim \text{Poisson}(\Lambda_i w_1)$  and  $M_i \sim \text{Poisson}(\Lambda_i w_2)$  for  $w_1 = 15, w_2 = 5$ .
- Claim sizes for each line are  $\text{Exp}(m_1 = 10)$  and  $\text{Exp}(m_2 = 40)$ .
- ABC posteriors based on 50  $X$ 's and on 250  $X$ 's with uniform priors.



Streftaris and Worton (2008), *Efficient and accurate approximate Bayesian inference with an application to insurance data*, Computational Statistics & Data Analysis

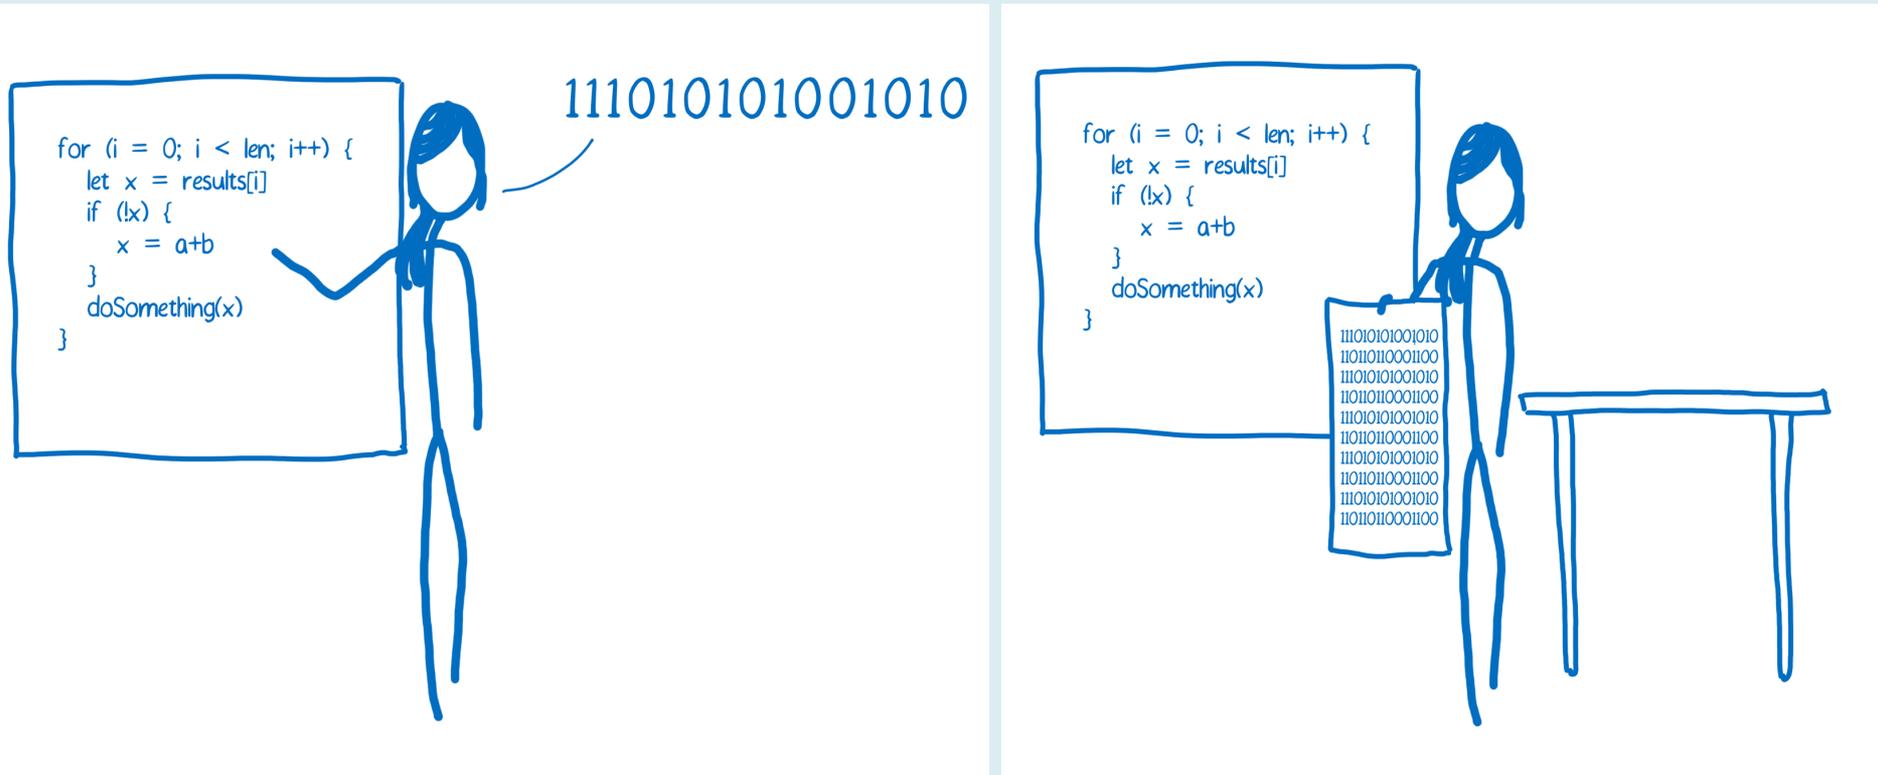
# ABC turns a statistics problem into a programming problem

... to be used as a last resort

# Compiled code

(use *numba*)

# Interpreters & compilers



C.f. Lin Clark (2017), *A crash course in JIT compilers*

## Interpreted version

```
1 def sample_geometric_exponential_sums(T, p, μ):
2     X = np.zeros(T)
3
4     N = rnd.geometric(p, size=T)
5     U = rnd.exponential(μ, size=N.sum())
6
7     i = 0
8     for t in range(T):
9         X[t] = U[i:i+N[t]].sum()
10        i += N[t]
11
12    return X
```

## Compiled by *numba*

```
1 from numba import njit
2
3 @njit()
4 def sample_geometric_exponential_sums(T, p,  $\mu$ ):
5     X = np.zeros(T)
6
7     N = rnd.geometric(p, size=T)
8     U = rnd.exponential( $\mu$ , size=N.sum())
9
10    i = 0
11    for t in range(T):
12        X[t] = U[i:i+N[t]].sum()
13        i += N[t]
14
15    return X
```

First run is compiling (500 ms), but after we are  
down from 2.7 ms to 164  $\mu$ s (16x speedup)

Original code: 1.7 s

Basic profiling with **snakeviz**: 5.5 ms, 310x speedup

+ Vectorisation/preallocation with **numpy**: 2.7 ms, 630x speedup

+ Compilation with **numba**: 164  $\mu$ s, 10,000x speedup

And potentially:

+ Parallel over 80 cores: say another 50x improvement,  
so overall 50,000x speedup.

# Take-home messages

- What is ABC?
- ABC turns a *statistics problem* into a *programming problem*
- `pip install approxbayescomp`