Neural networks and uncertainty

Actuarial AI Day

Patrick Laub, UNSW

Joint work with Benjamin Avanzi, **Eric Dong**, and Bernard Wong

Car insurance

Claim size prediction

👤 Age	🚗 Age	🏎️ Type
25	3	🚙 Sedan
40	5	🚐 SUV
19	1	🏎️ Sports Car
60	10	🚗 Hatchback

Car insurance

Claim size prediction

👤 Age	🚗 Age	🏎️ Type
25	3	🚙 Sedan
40	5	🚐 SUV
19	1	🏎️ Sports Car
60	10	🚗 Hatchback

Car insurance

Claim size prediction

👤 Age	🚗 Age	🏎️ Type		Cost
25	3	🚙 Sedan		💵 $1,200
40	5	🚐 SUV	⟶	💵 $2,500
19	1	🏎️ Sports Car		💵 $3,800
60	10	🚘 Hatchback		💵 $800

Car insurance

Claim size prediction

👤 Age	🚗 Age	🏎️ Type		Cost
25	3	🚙 Sedan		💵 $1,200
40	5	🚐 SUV	⟶	💵 $2,500
19	1	🏎️ Sports Car		💵 $3,800
60	10	🚗 Hatchback		💵 $800

Deterministic model! Move to a probabilistic model.

Distributional regression

Customer 1 = (25, 3, 🚙)

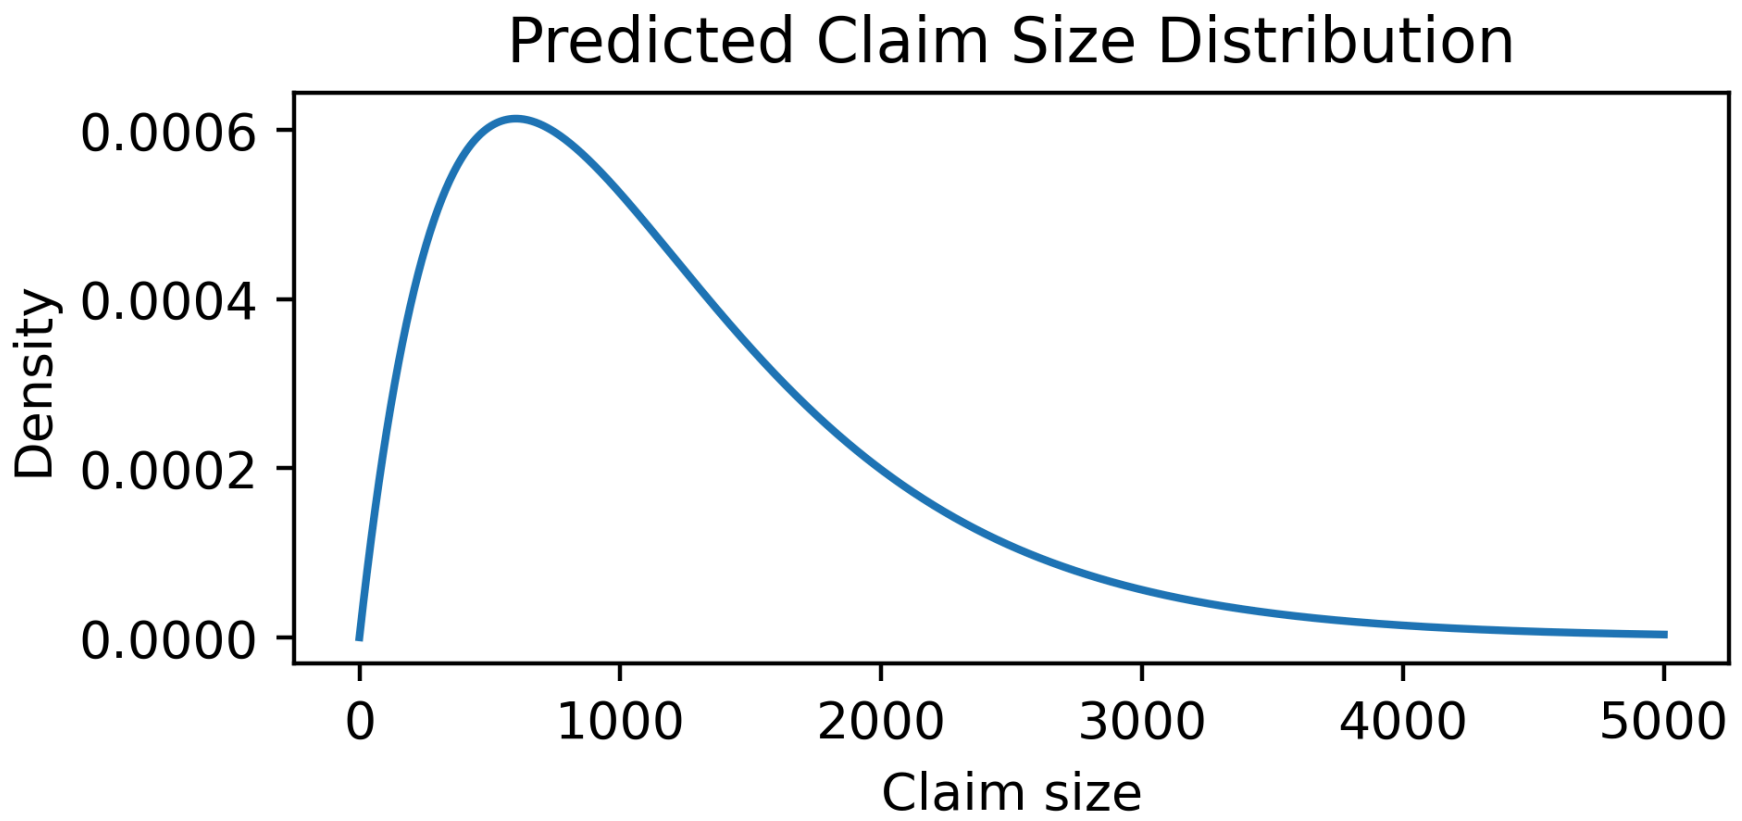

Distributional regression

Customer 2 = (40, 5, 🚐)

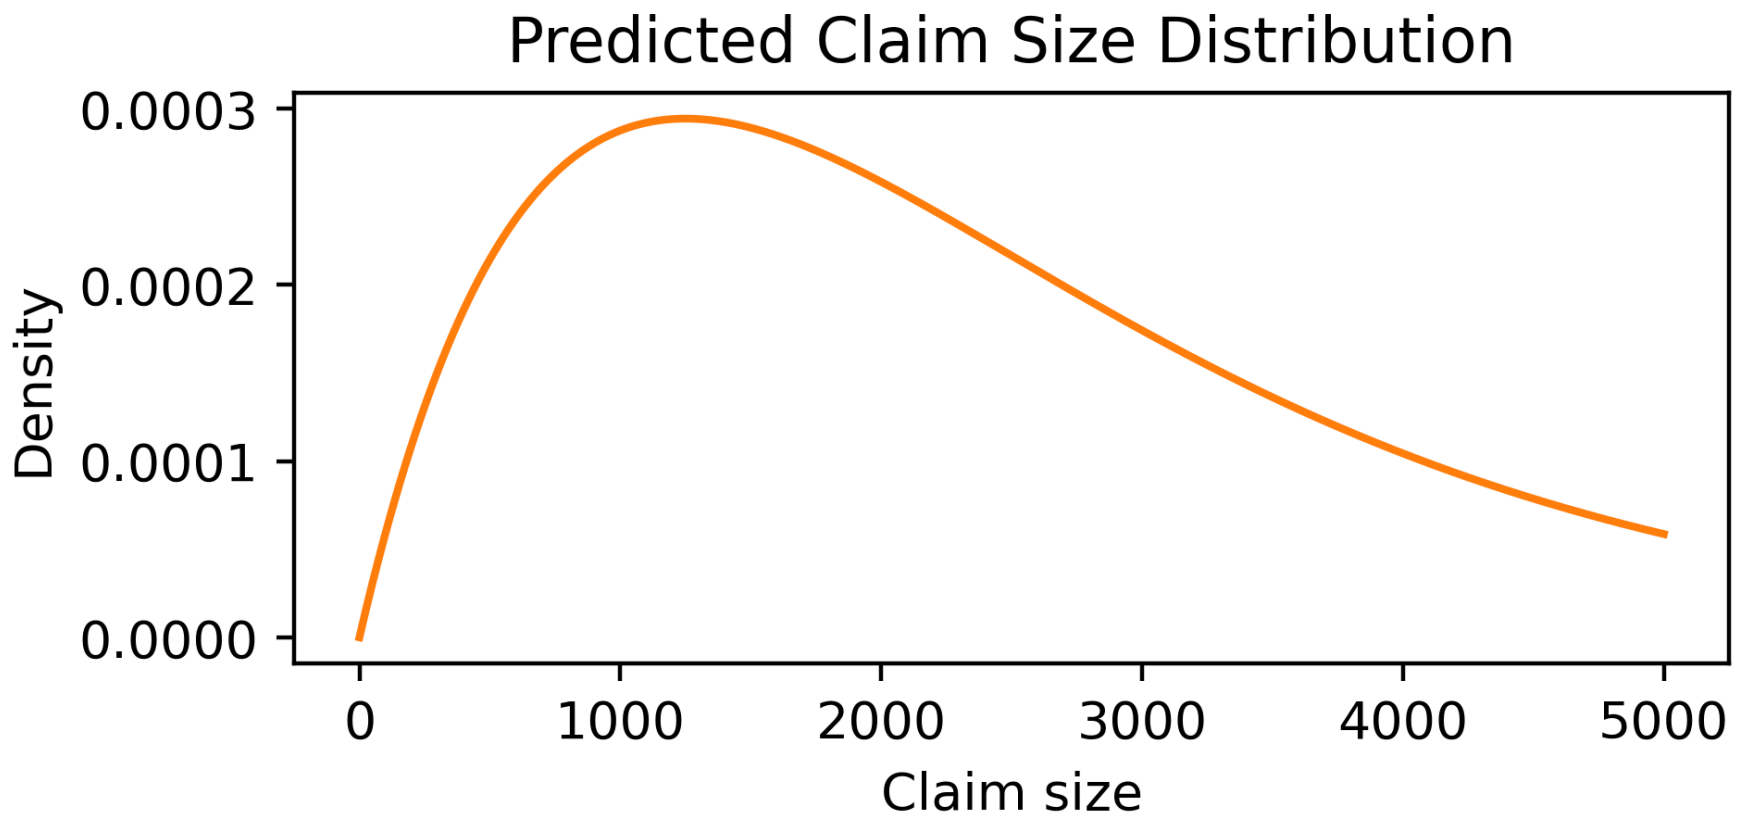

Distributional regression

Customer 3 = (19, 1, 🏎️)

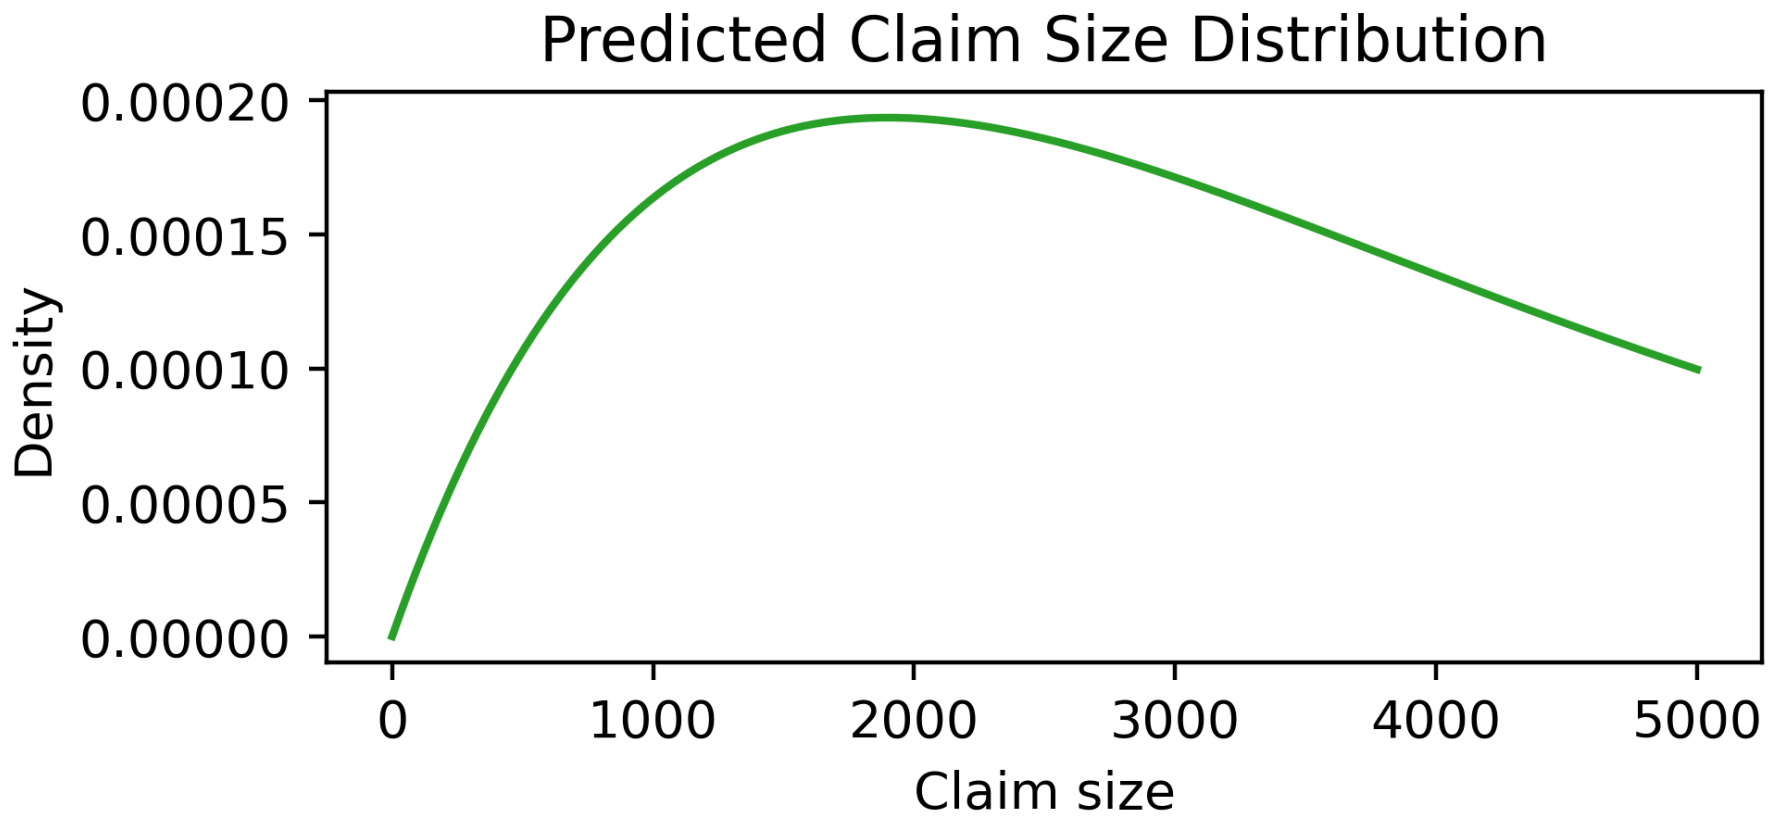

Distributional regression

Customer 4 = (60, 10, 🚗)

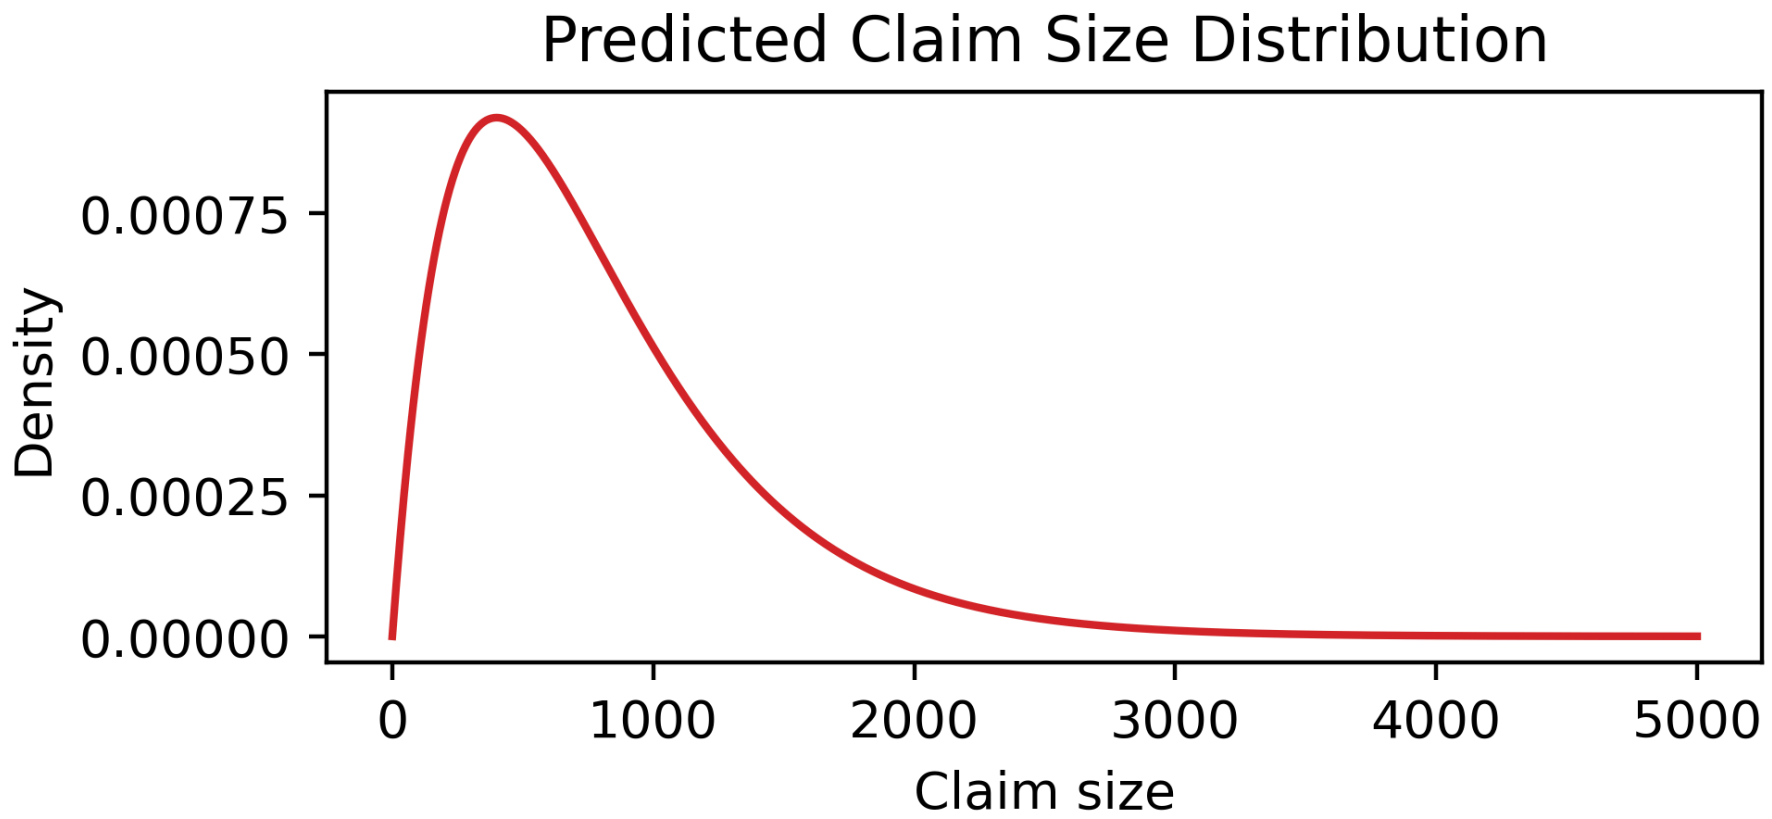

Distributional regression

All customers

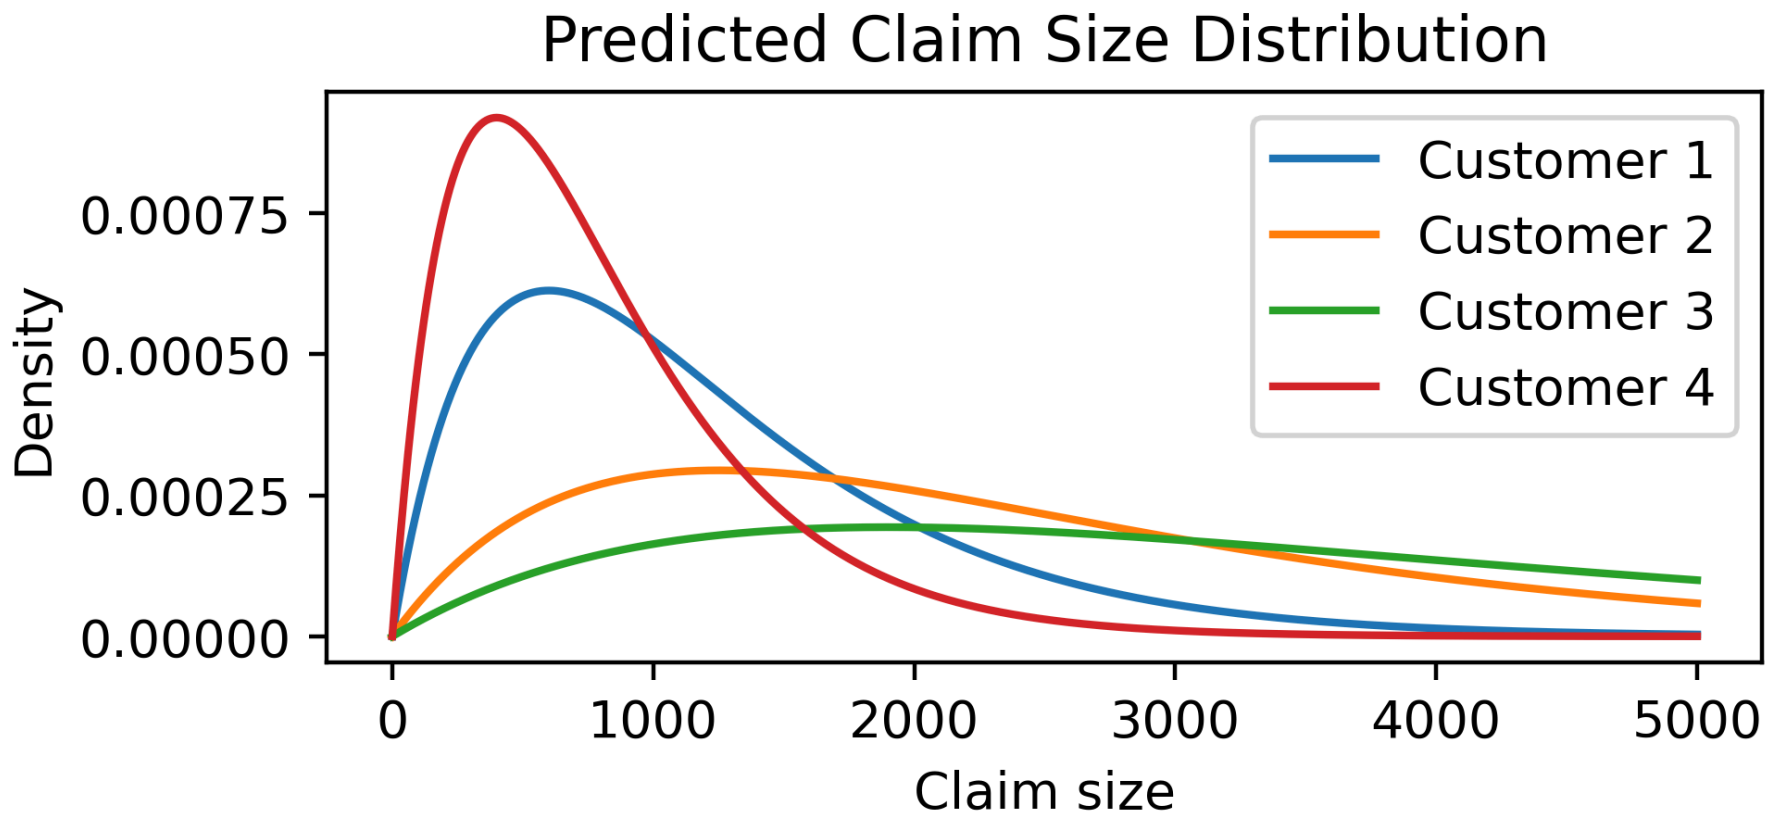

A generalised linear model

A gamma GLM with a log link function:

$$Y|\mathbf{X} \sim \text{Gamma}(\ldots, \ldots)$$

$$\mathbb{E}[Y|\mathbf{X}] = \exp\Big\{\beta_0 + \beta_1 \cdot \text{Age} + \beta_2 \cdot \text{Car Age} + \beta_3 \cdot \text{Type}\Big\}$$

A simple model, easy to train and interpret, but…

A generalised linear model

A gamma GLM with a log link function:

$$Y|\mathbf{X} \sim \text{Gamma}(\dots, \dots)$$

$$\mathbb{E}[Y|\mathbf{X}] = \exp\Big\{\beta_0 + \beta_1 \cdot \text{Age} + \beta_2 \cdot \text{Car Age} + \beta_3 \cdot \text{Type}\Big\}$$

A simple model, easy to train and interpret, but…

GLMs can be

1. Bad at *regression*
2. Bad at *distributional* regression

Example 1: Non-monotonicity

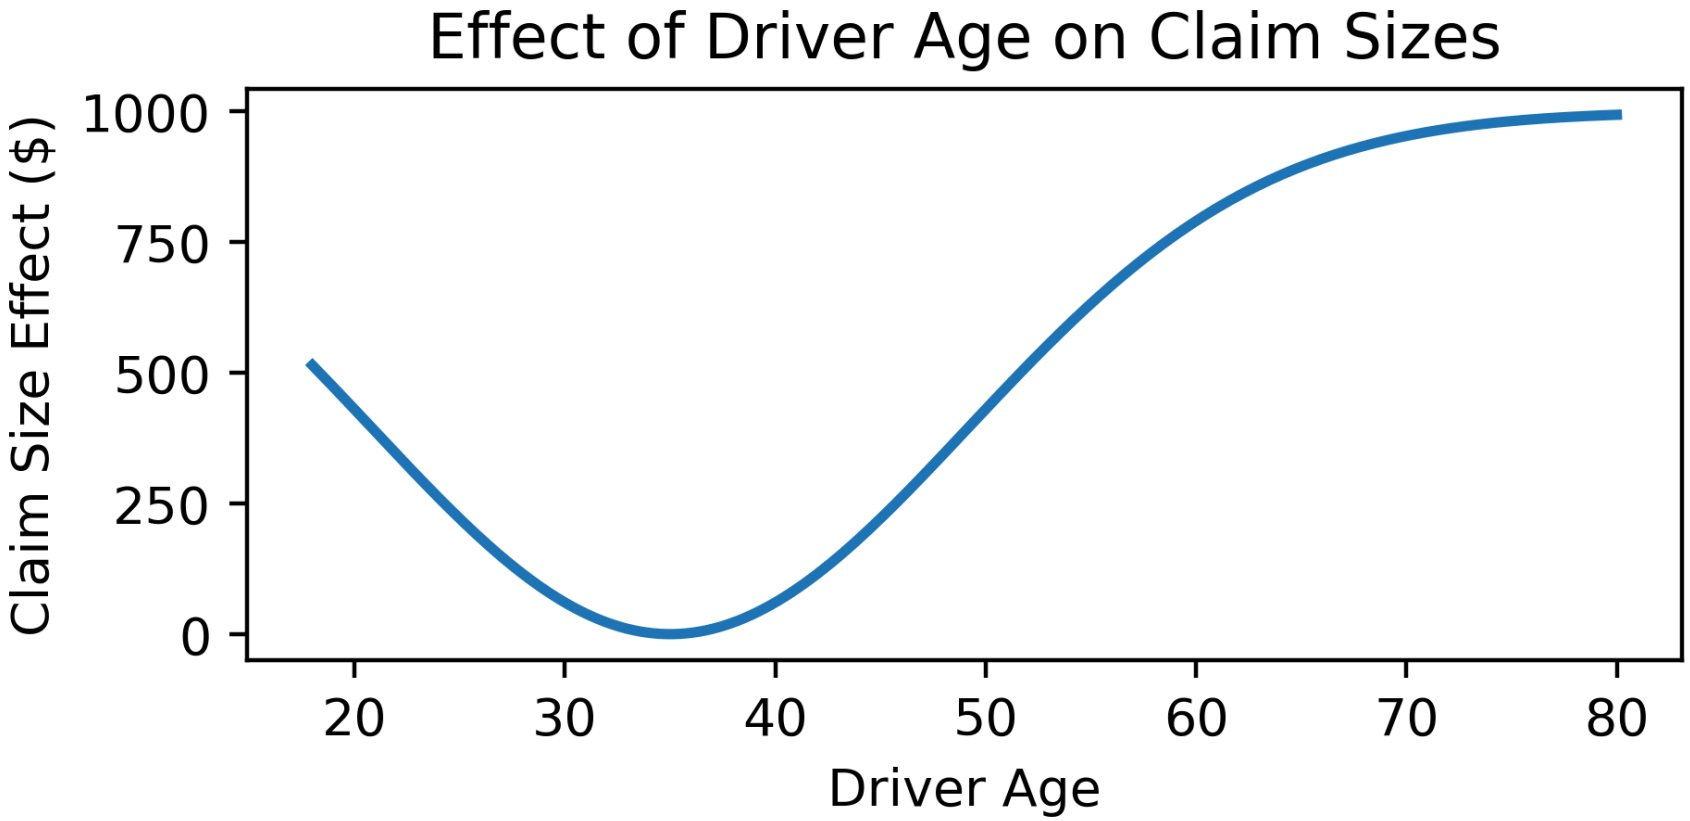

GLMs cannot (easily) do this $\longrightarrow$ Use a neural network

Example 2: Multi-modality

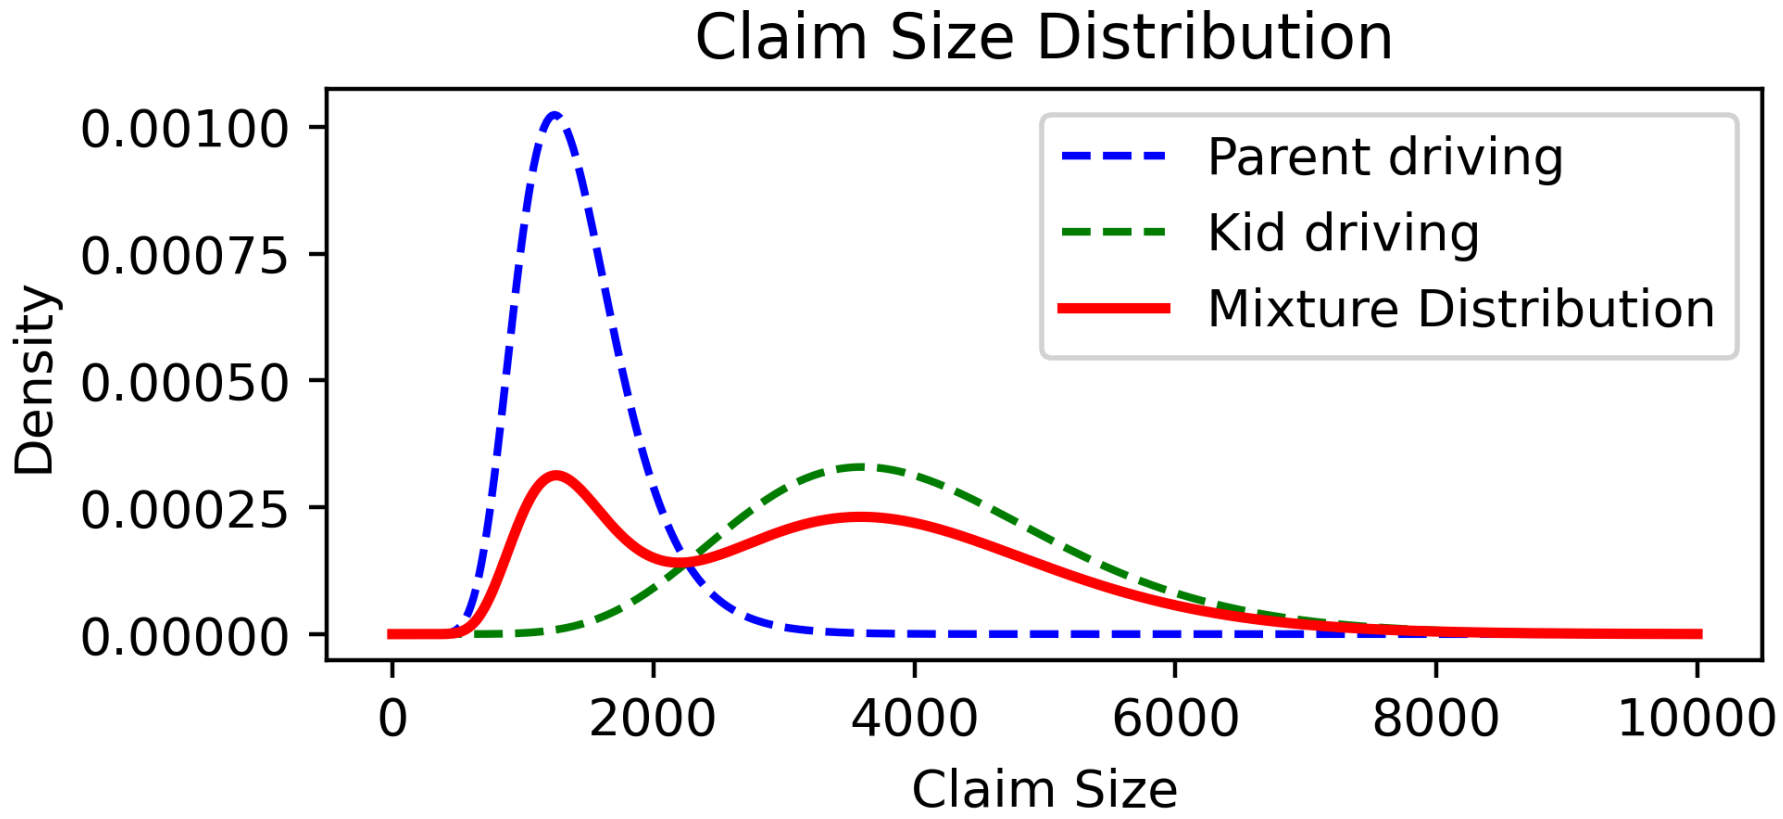

GLM distributions are too restrictive $\longrightarrow$ Adjust them

Current solutions

- The Generalized Linear Model (GLM; Nelder and Wedderburn, 1972) predicts the mean assuming exponential family distribution.
- The Combined Actuarial Neural Network (CANN; Schelldorfer and Wüthrich, 2019) adjusts the mean predictions of the GLM through a neural network.
- The Mixture Density Network (MDN; Bishop, 1994) predicts the distributional parameters of a mixture distribution.
- The Deep Distribution Regression (DDR; Li et al., 2021) directly models the distribution as a mixture of uniform distributions.

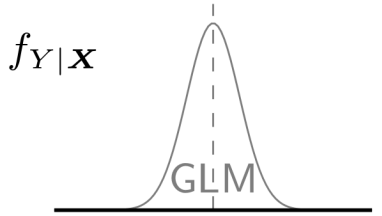

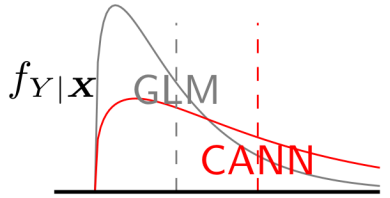

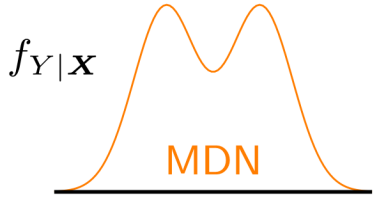

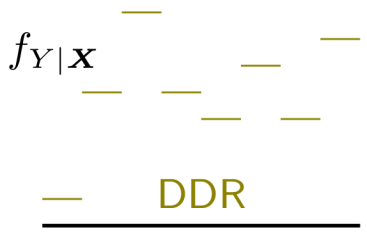

Thanks to Eric Dong!

UNSW SYDNEY

Distributional Refinement Network

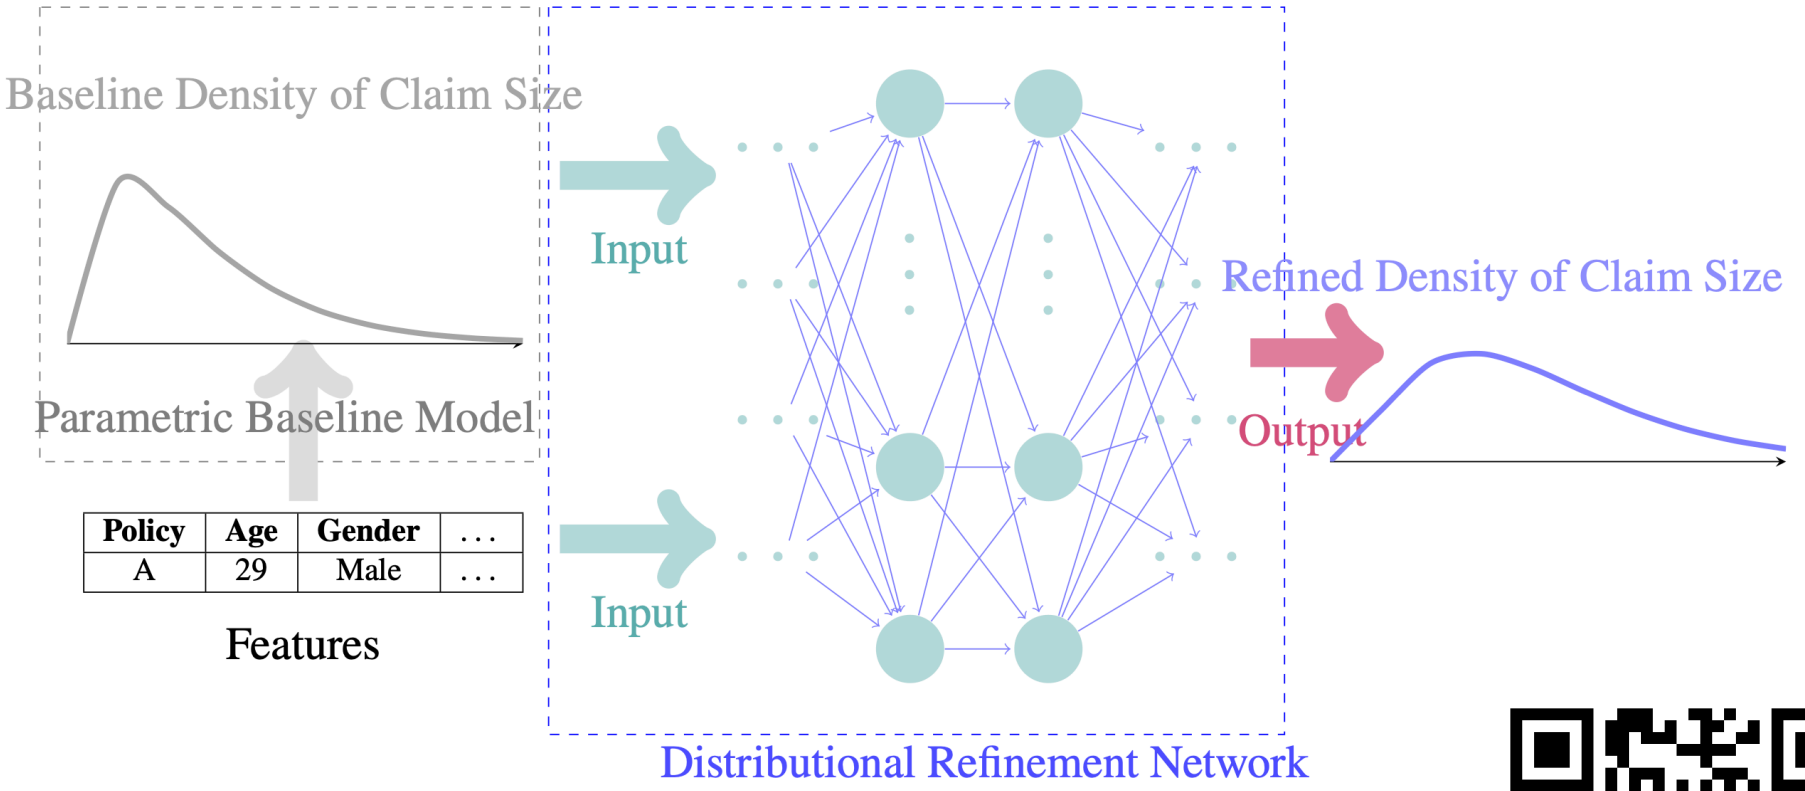

Use a network to *adjust* the distributions out of a GLM.
The user can select small or large changes (botox or a facelift!).
Check out the paper and the 'drn' Python package!

Joint work by Benjamin Avanzi, **Eric Dong**, Patrick Laub, and Bernard Wong.