Actuaries Institute.

Optimal Illusion – A look at the practical limitations of price optimisation

Hugh Miller and Tom Moulder

May 2024

Presented at the 2024 All Actuaries Summit

The Actuaries Institute acknowledges the traditional custodians of the lands and waters where we live and work, travel, and trade.

We pay our respect to the members of those communities, Elders past and present, and recognise and celebrate their continuing custodianship and culture.

Actuaries Institute.

Introduction

- What is price optimisation?
- Why is price optimisation a big deal in insurance (especially personal lines in general insurance)?
- Why don't we talk about optimisation much?
- What happens in other jurisdictions?

Actuaries Institute.

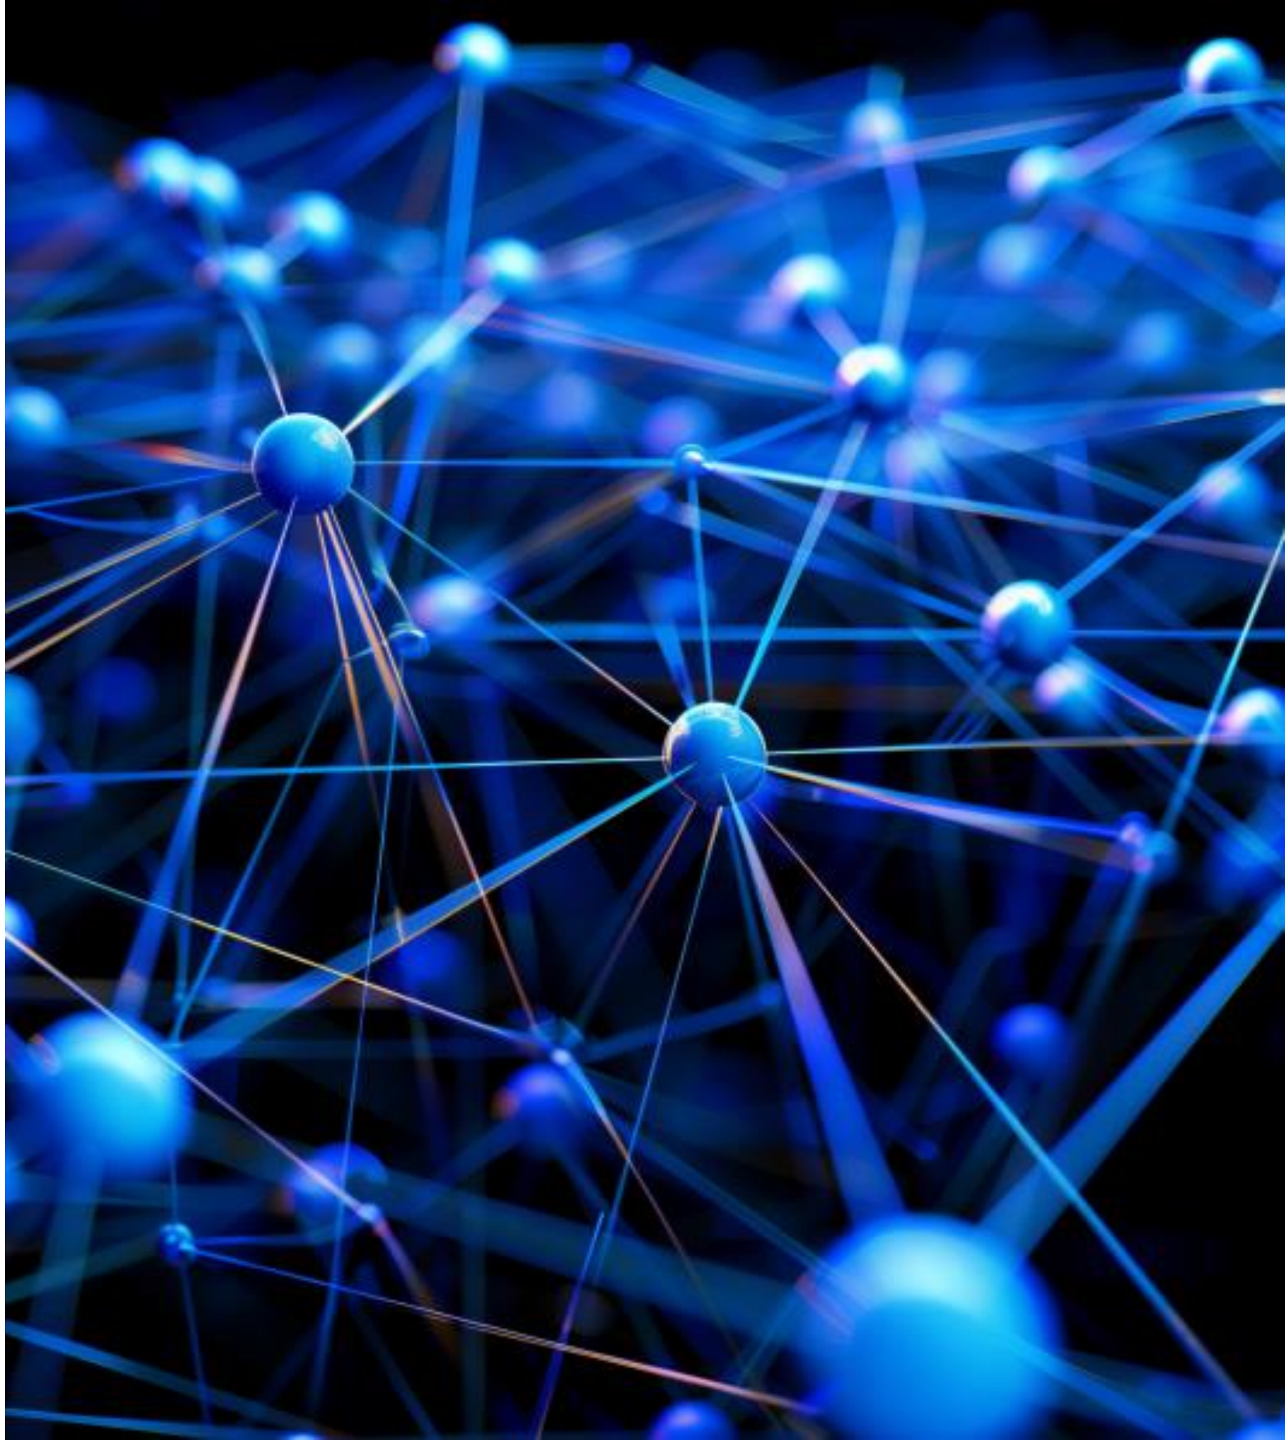

Agenda

An overview of our paper (**more detail there!**) :

1. Simulated data setup

2. Issues to consider when applying optimisation
   - Choice of constraints in optimisation
   - Drivers of price differentiation
   - Issues with overconfidence in elasticity estimates
   - Issues with model specification
   - Multiyear considerations
   - Estimation and validation of demand & elasticity

3. Concluding points

Actuaries Institute.

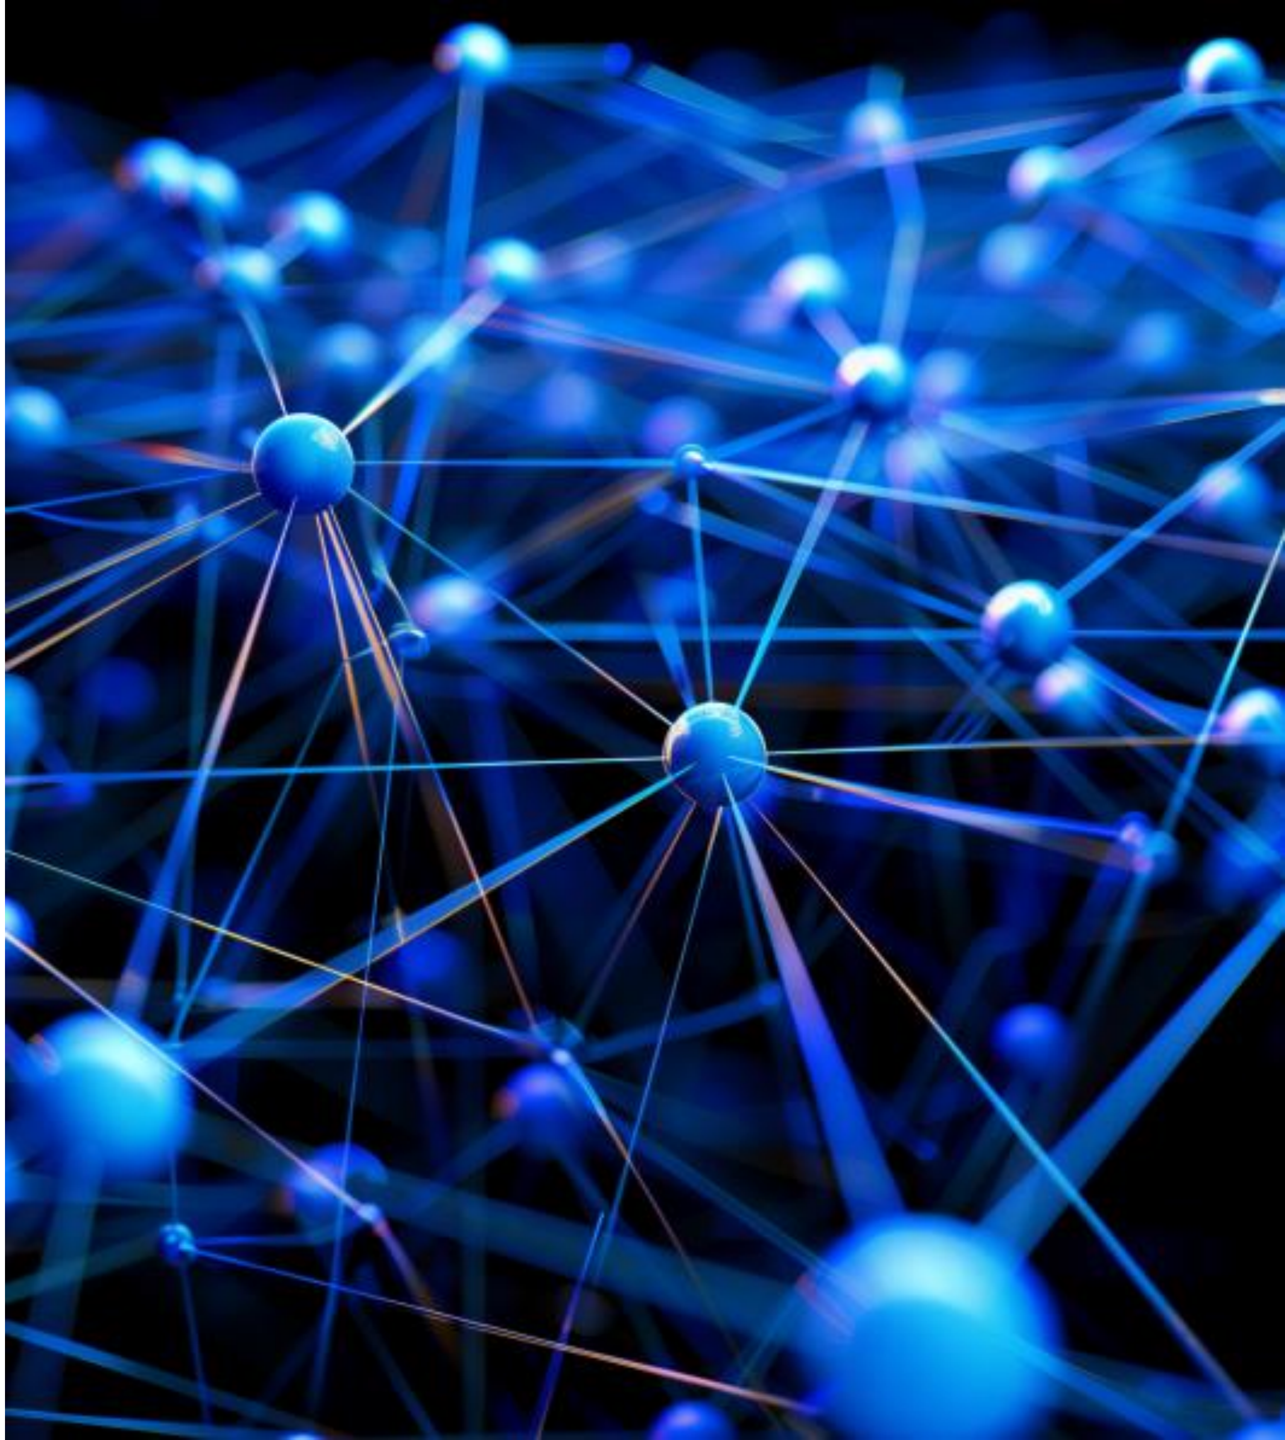

Simulated data setup

01

Actuaries Institute.

Setup – simplifications

We simulate an insurer's renewal book. Designed to be plausible with flexibility, albeit still several simplifications:

- Simple relationship between initial profit margins and risk (% loading)
- Fixed point in time optimisation (all renewing simultaneously)
- Focus on renewals only (as opposed to new business)
- Single line of business
- No competitor premium effects
- Smooth, convex demand vs price relationships

All the above solvable, but of lesser interest for our paper

Setup – technical

Assumptions

- Starting profit margin (before optimisation) of 15%
- Cup and cap on price changes of ±20%
- Overall elasticity low-ish (global price rise increases revenue)
- Average retention rate of 80%
- $n = 100{,}000$ customers, indexed by i
- Normal IID latent variables x_{i1}, x_{i2}, x_{i3} and x_{i4} driving the variability and correlations between retention, risk and elasticity.

Optimiser

Built bespoke optimiser routine to target a specific volume of risk cost (sum of total risk cost retained) and maximise profit subject to the constraint. Leverages derivative calculations of profit against risk cost, $\partial\pi_i/\partial r_i$, and runs quickly (<3 seconds).

Parametrisation

A retention model using inverse logistic link $g(x)$. Demand is

$$d(x_i) = g(g^{-1}(\mu_d) + \phi_1 x_{i1} + x_{i3}) = g(z_i)$$

Demand then modified by a price change p_i and elasticity e_i,

$$d(x_i, p_i) = g(z_i + e_i p_i)$$

A risk cost model R_i with log link for the mean

$$E(R_i) = \exp(ln(800) + x_{i1} + x_{i2})$$

Risk is then sampled from a gamma distribution with parameters drive by sampled means and a constant coefficient of variation $\frac{\sigma_{ri}}{\mu_{ri}} = 0.4$

Profit as a function of premium and risk cost $\pi_i = P_i - R_i$.

An elasticity parameter (subject to a maximum of −0.5) of

$$e_i = -\left(3 + 2(\phi_2 x_{i1} + x_{i4})\right)$$

Actuaries Institute.

Setup – technical (2)

Combining the above, expected profit can be written down, for customer i, as

$$\pi_i = \left(P_i(1 + p_i) - R_i\right) g(z_i + e_i p_i)$$

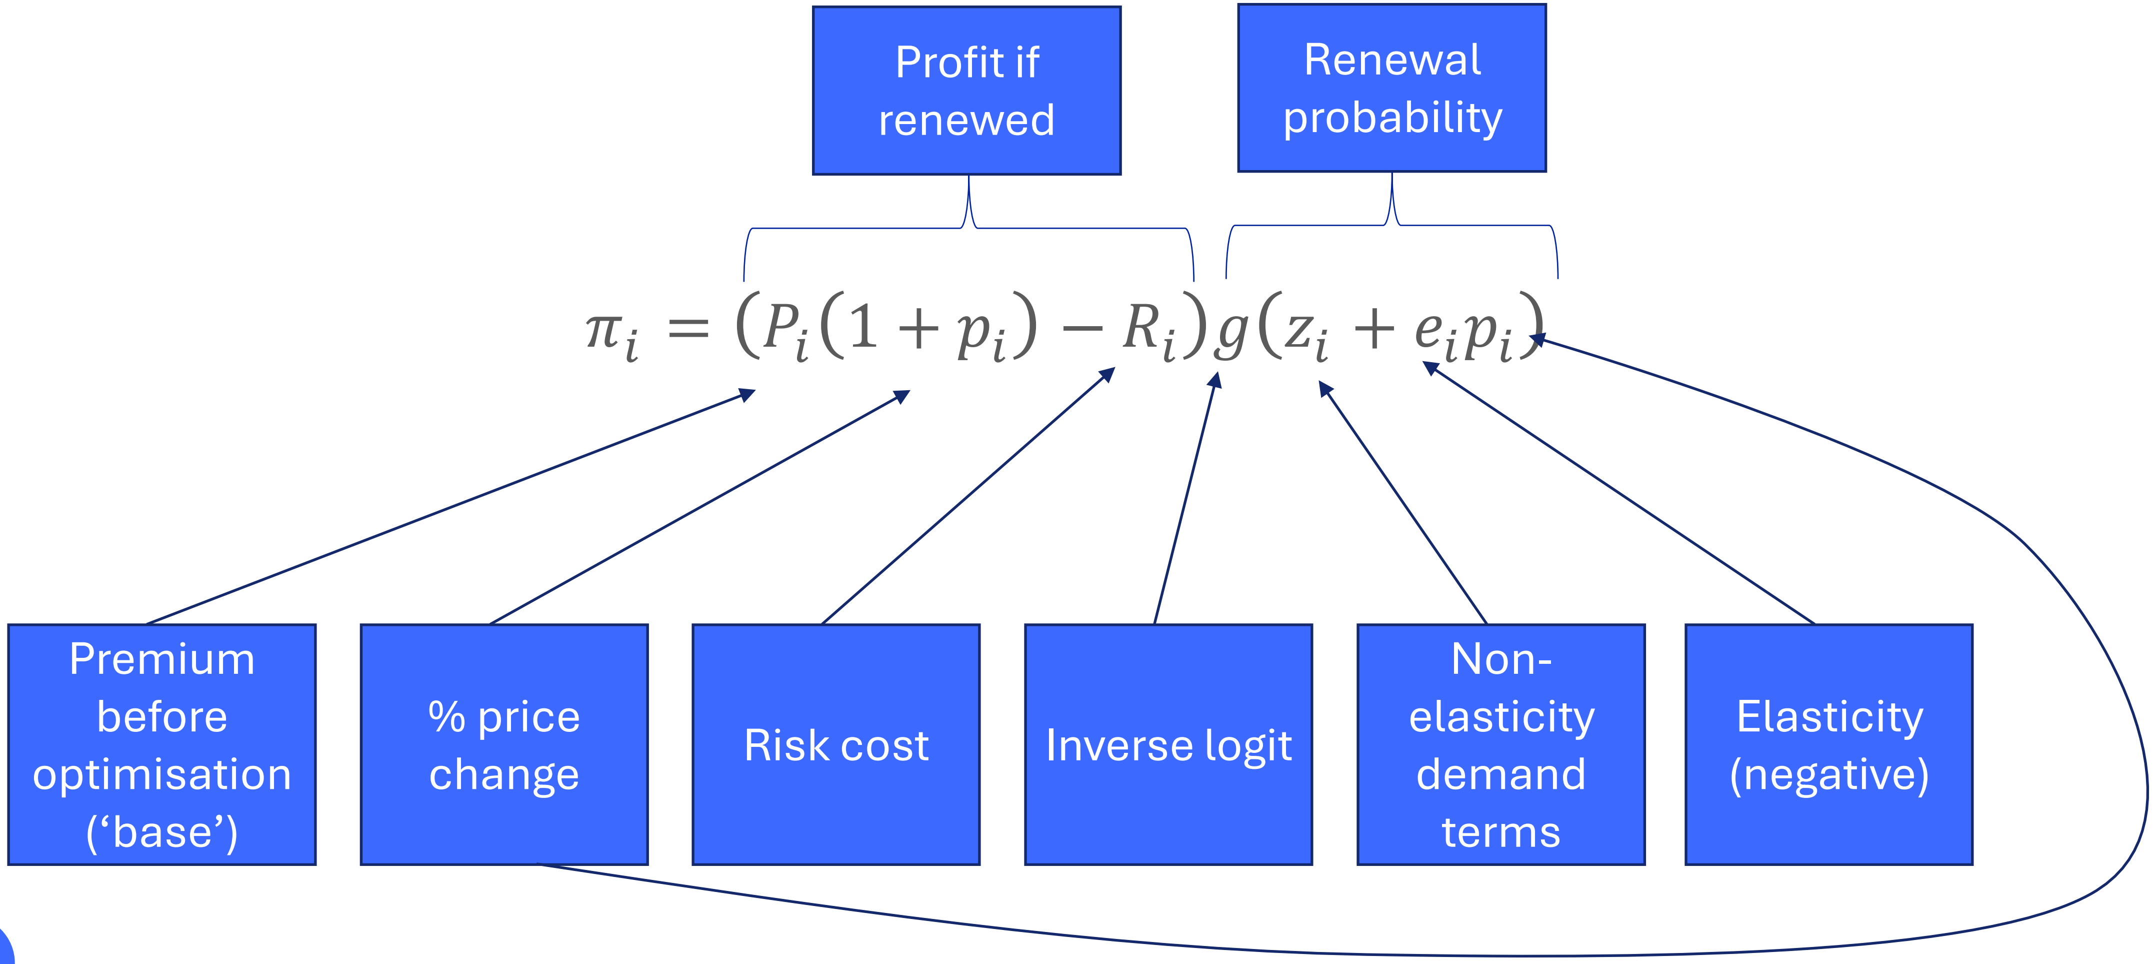

Actuaries Institute.

Basic optimisation results

Targeting a risk volume close to neutral for our dataset ($68m), we see a 19% increase in profit, equalling about 3% of risk cost.

Overall gains in this setup are small to moderate when compared to gains achieved by headline pricing changes

Scenario	Cust volume	Avg prem	GWP ($m)	Risk prem ($m)	Profit ($m)
All price change 0%	**79,578**	**1,004.49**	**79.98**	**67.98**	**12.00**
Uniform price change	79,600	1,004.04	79.96	68.00	11.96
Optimised prices	79,712	1,036.20	82.25	68.00	14.25

Standard and optimised curves trading off risk volumes and profit

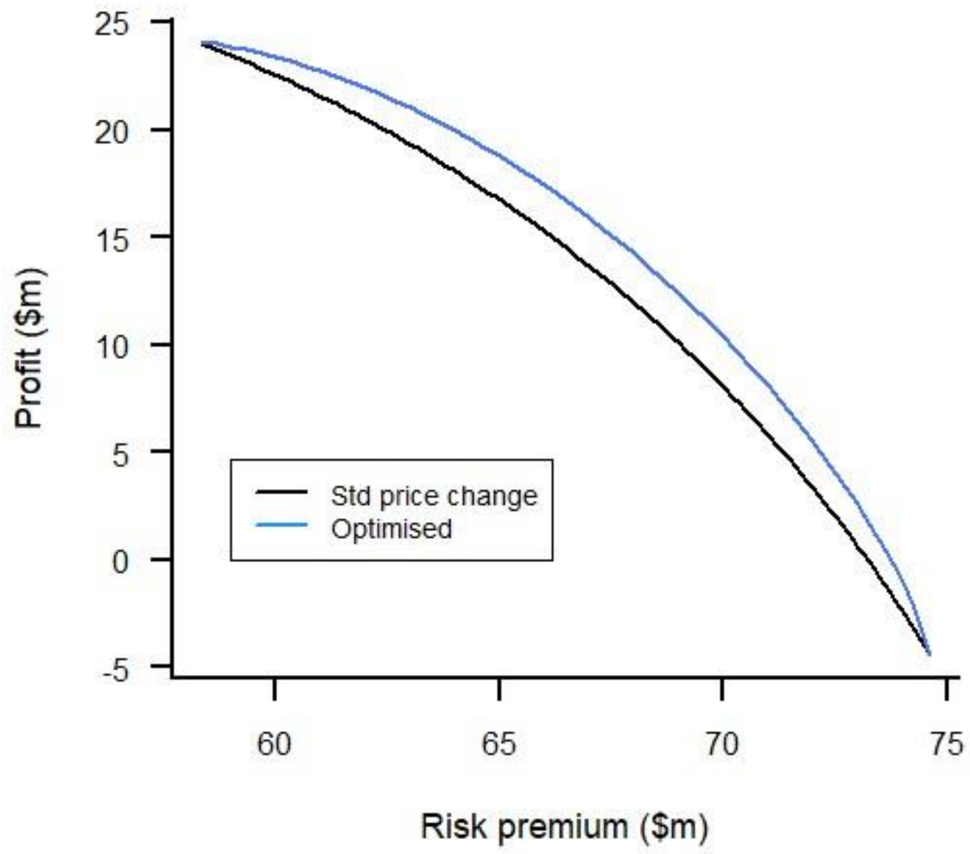

Actuaries Institute.

Basic results – basic optimisation (2)

1 *Price change distribution*

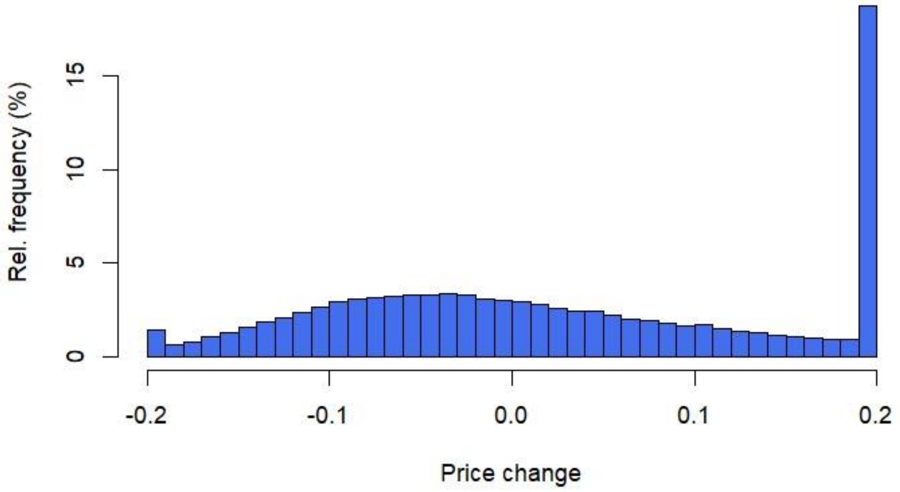

2 *Price change by initial retention rate*

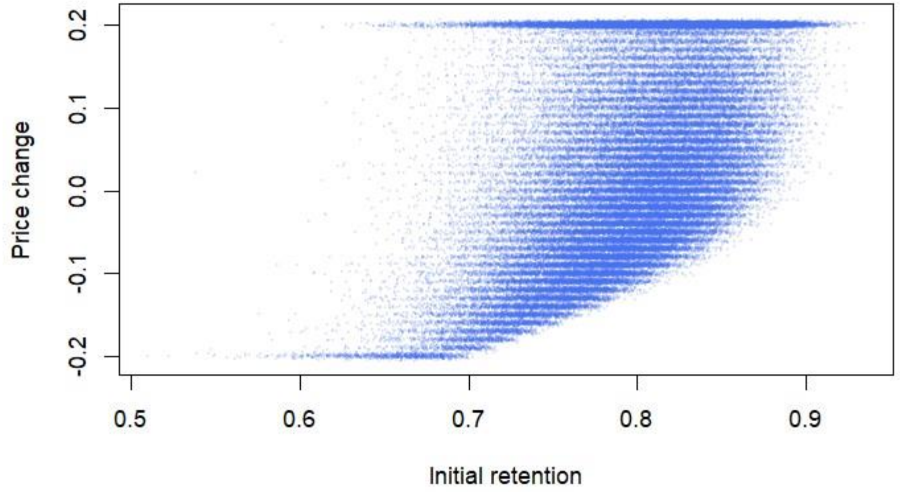

3 *Price change by elasticity parameter*

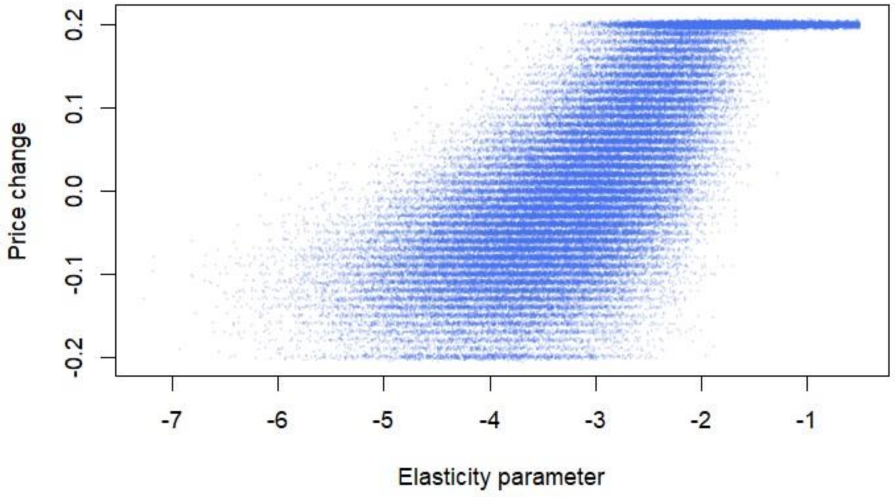

4 *Price change by base premium*

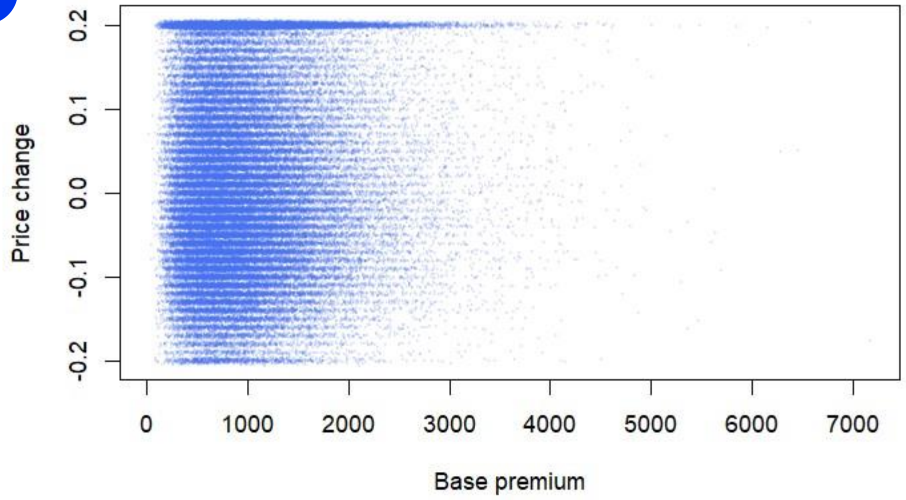

Issues to consider when applying optimisation

02

Actuaries Institute.

02 – Issues to consider when applying optimisation

A. **Choice of constraints in optimisation**
B. Drivers of optimised price changes
C. Overconfidence in elasticity estimates
D. Issues with model misspecification
E. Impact on the customer base over time
F. Challenges of modelling elasticity

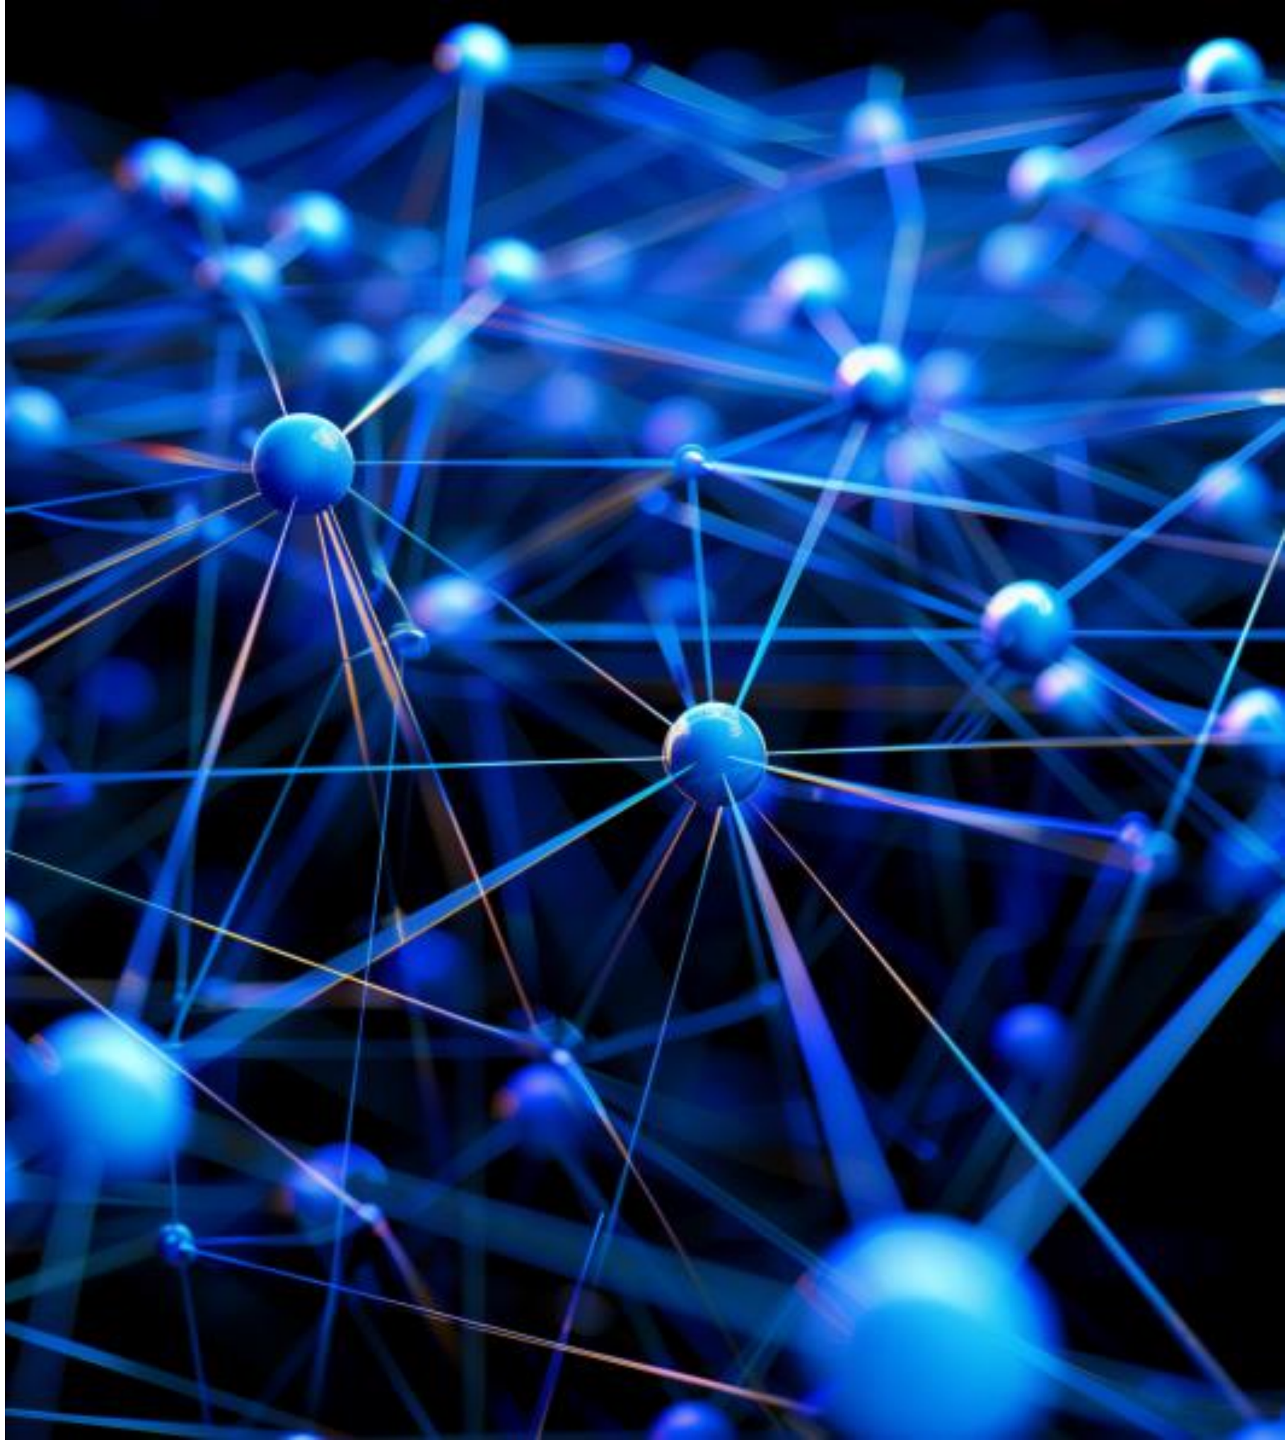

Actuaries Institute.

Choice of constraints in optimisation

Based on our reading, it is not universal to use risk premium volume as the primary constraint.

Another natural choice is the **absolute number of customers**. That is, maximise profits for a given number of customers retained.

At face value, this generates (in our opinion) perverse impacts by maximising price increases for larger policies and reductions for smaller ones. An example is shown right.

Retaining small premium customers can lead to:

- Reductions in GWP over time.
- A rising expense ratio (if expenses include elements of fixed cost per policy)

Price changes by starting premium obtained when optimising for customer numbers rather than risk premium volumes

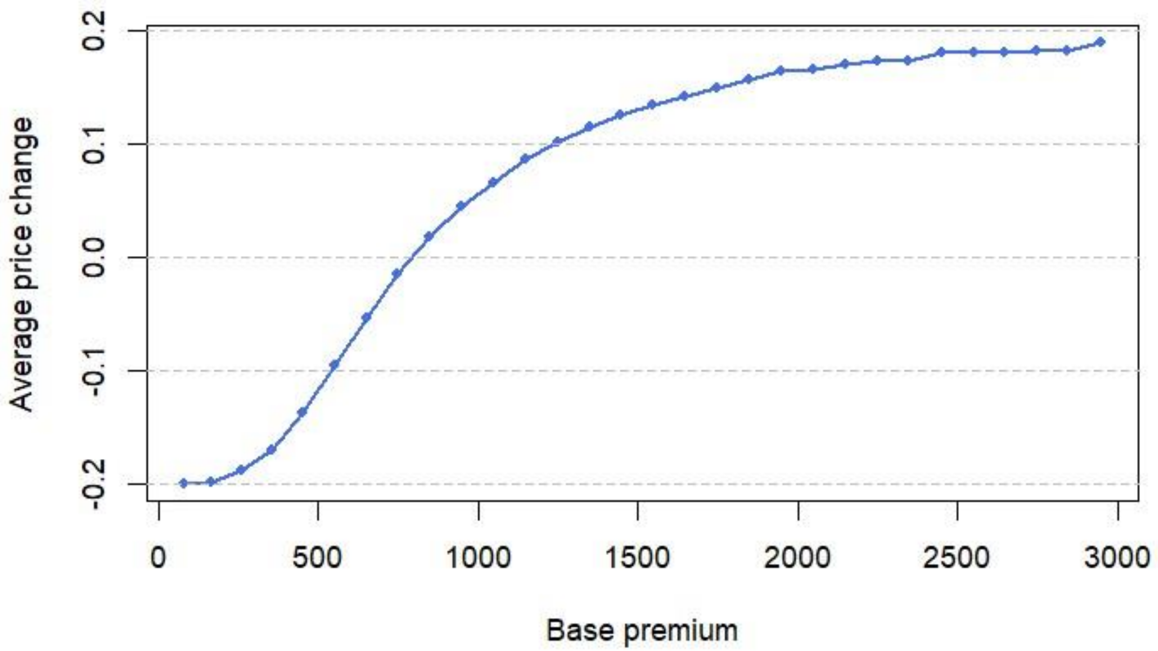

Poorly chosen constraints will cause unusual skews in the customer base

Actuaries Institute.

02 – Issues to consider when applying optimisation

A. Choice of constraints in optimisation
B. Drivers of optimised price changes
C. Overconfidence in elasticity estimates
D. Issues with model misspecification
E. Impact on the customer base over time
F. Challenges of modelling elasticity

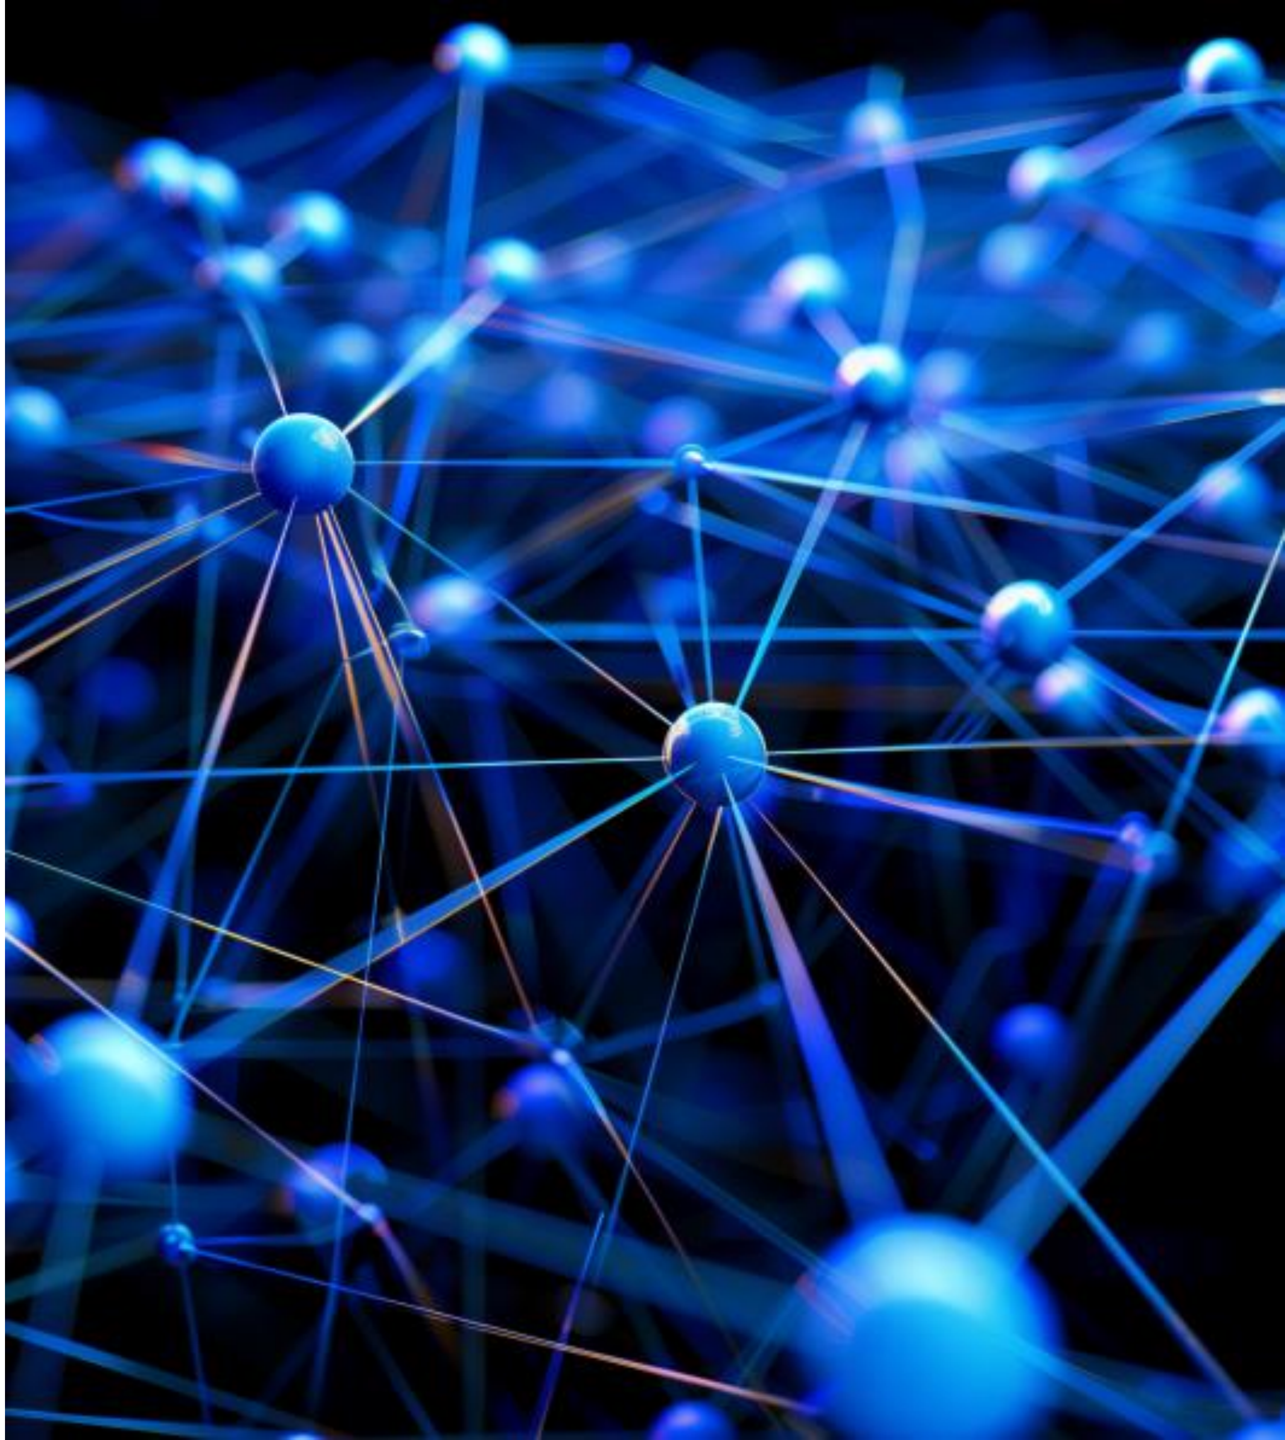

Actuaries Institute.

Drivers of optimised price changes

Under a logit-link model, the key derivative (profit per unit demand) is:

$$\frac{\partial \pi}{\partial r} = \frac{P(1+p)-R}{R} + \frac{P}{eR\left(1-d(x)\right)}$$

Elasticity parameter e and base demand $d(x)$ both appear. Assuming that e is always negative and we want to increase prices when $\partial\pi/\partial r$ is low:

- More negative e (more elastic) → higher $\partial\pi/\partial r$ → premium decrease
- Higher base demand $d(x)$ → lower $\partial\pi/\partial r$ → premium increase

Since both demand and retention affect the optimisation, their relative level and variability influence how much they contribute to the result.

Approximate proportion of variation in $\partial\pi/\partial r$ due to variation in initial retention

Data with average initial retention of 80%.

Standard deviation of elasticity e_i	Standard deviation of initial retention parameter z_i						
	0	0.1	0.2	0.3	0.4	0.5	0.6
0.0		100%	100%	100%	100%	100%	100%
0.2	0%	54%	71%	79%	84%	88%	90%
0.4	0%	36%	53%	64%	71%	77%	81%
0.6	0%	25%	40%	51%	59%	66%	72%
0.8	0%	16%	28%	38%	46%	53%	60%
1.0	0%	11%	20%	28%	35%	42%	49%
1.2	0%	8%	15%	21%	28%	34%	40%
1.4	0%	6%	12%	18%	23%	29%	35%
1.6	0%	5%	11%	16%	21%	26%	31%

Optimisation driven by elasticity

Optimisation driven by retention

Both elasticity and base demand impact optimised premiums.
Choice of link will affect this relationship

Actuaries Institute.

02 – Issues to consider when applying optimisation

A. Choice of constraints in optimisation

B. Drivers of optimised price changes

C. Overconfidence in elasticity estimates

D. Issues with model misspecification

E. Impact on the customer base over time

F. Challenges of modelling elasticity

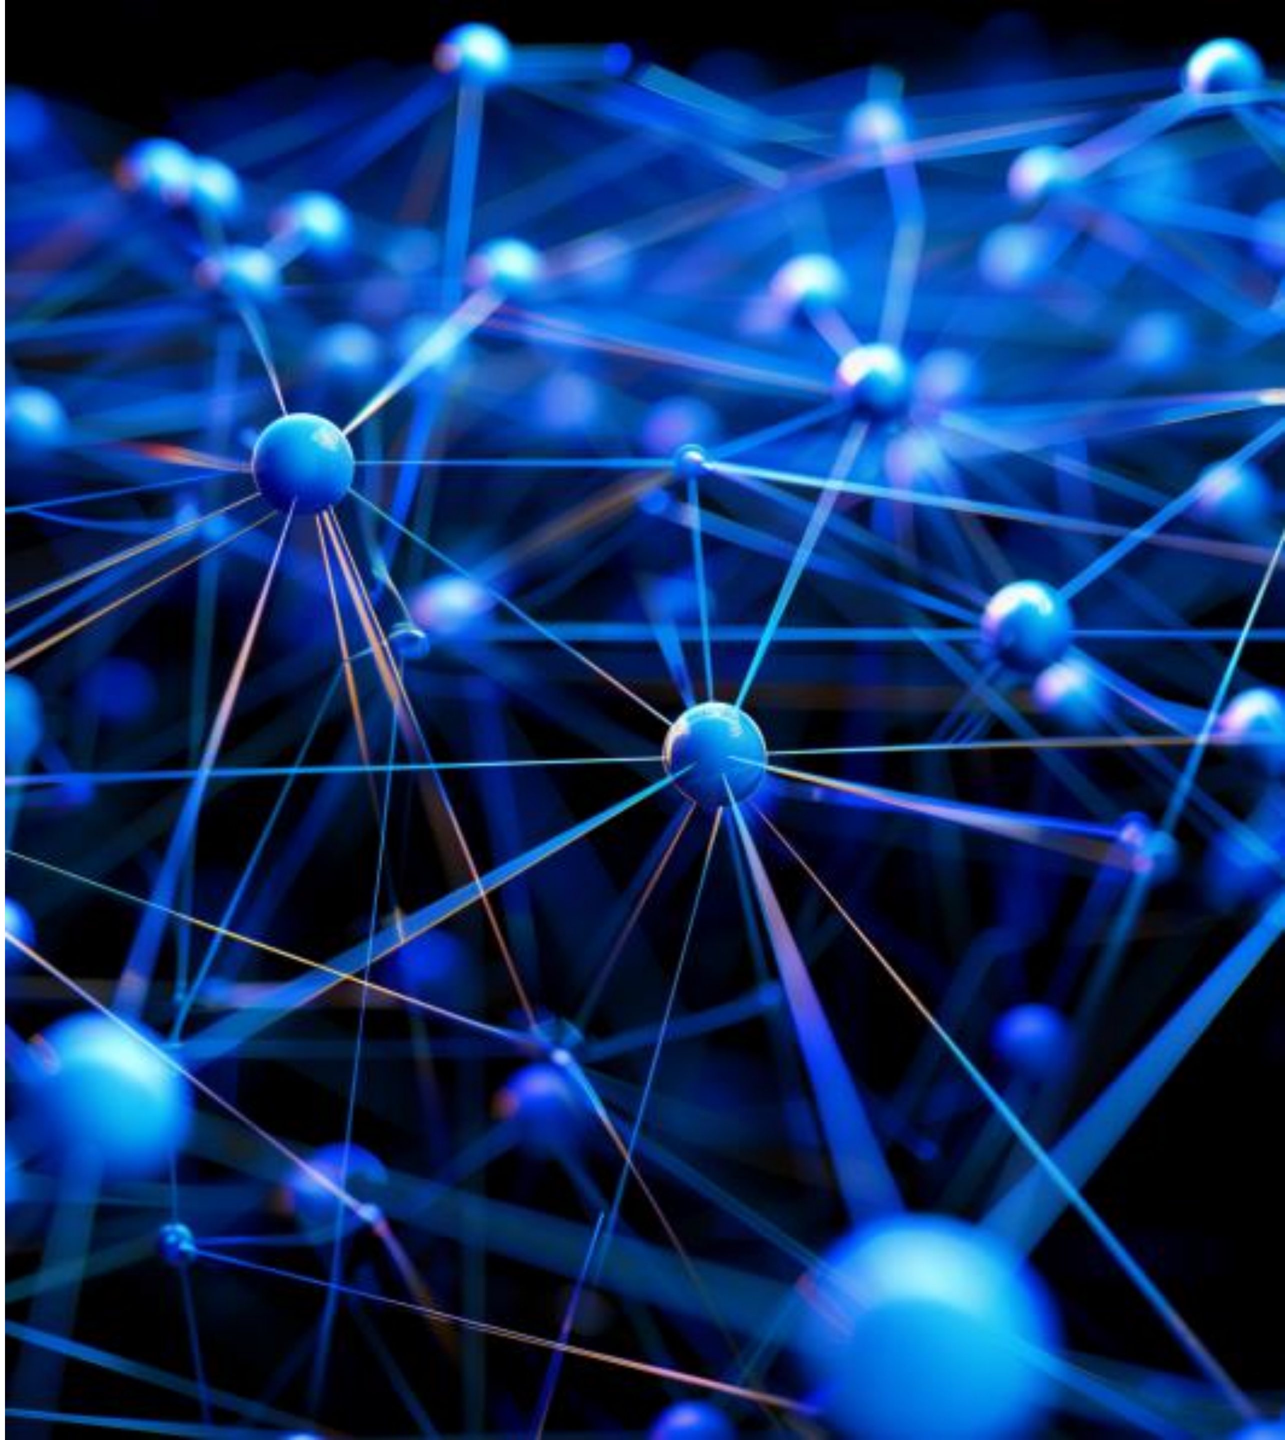

Actuaries Institute.

Overconfidence in elasticity estimates

Semenovich & Petterson (2019) showed, under simplified small-variation assumptions, that undiagnosed uncertainty in the elasticity estimate significantly degrades optimisation performance.

If σ_e is the std deviation of elasticity e and σ_a is the error in its estimation, we define T the correlation between true and estimated elasticity

$$T = \frac{\sigma_e}{\sqrt{\sigma_a^2 + \sigma_e^2}}$$

The existing result suggested actual ÷ expected profit boost equals a variant of T.

We can explore this result in a full simulated environment, holding the distributions of actual and expected elasticity constant.

Comparison of estimated and true elasticity parameter for different values of accuracy T

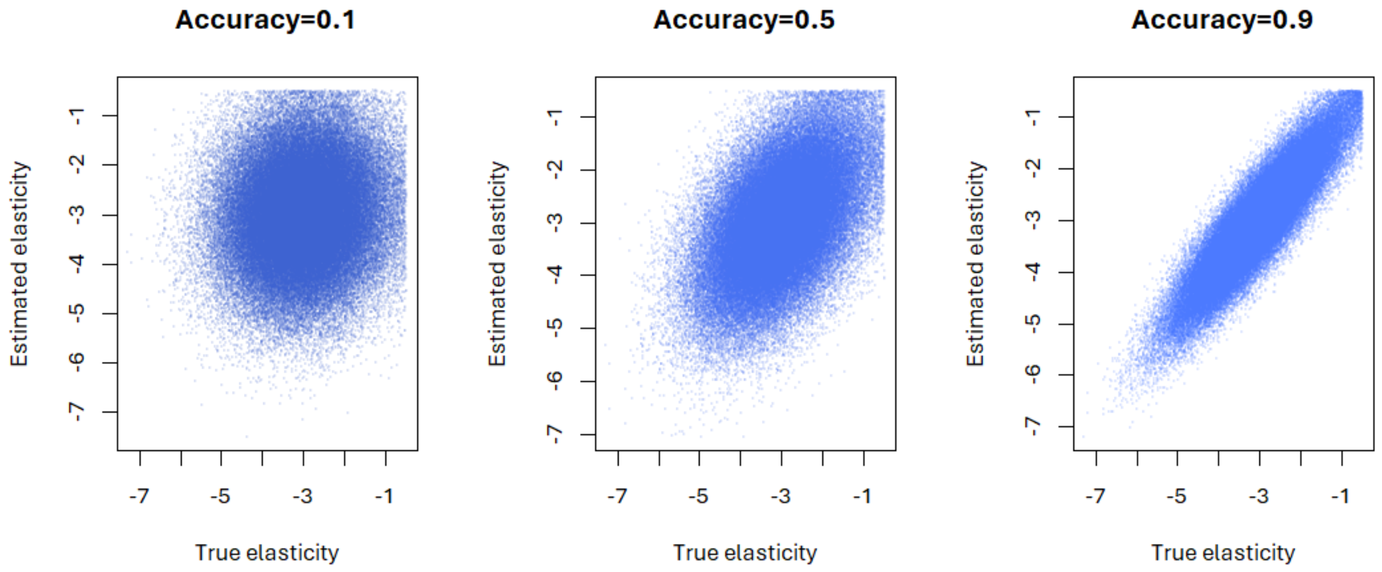

Actuaries Institute.

Overconfidence in elasticity estimates (2)

Expected and actual optimisation performance compared to standard price changes, $T = 0.5$

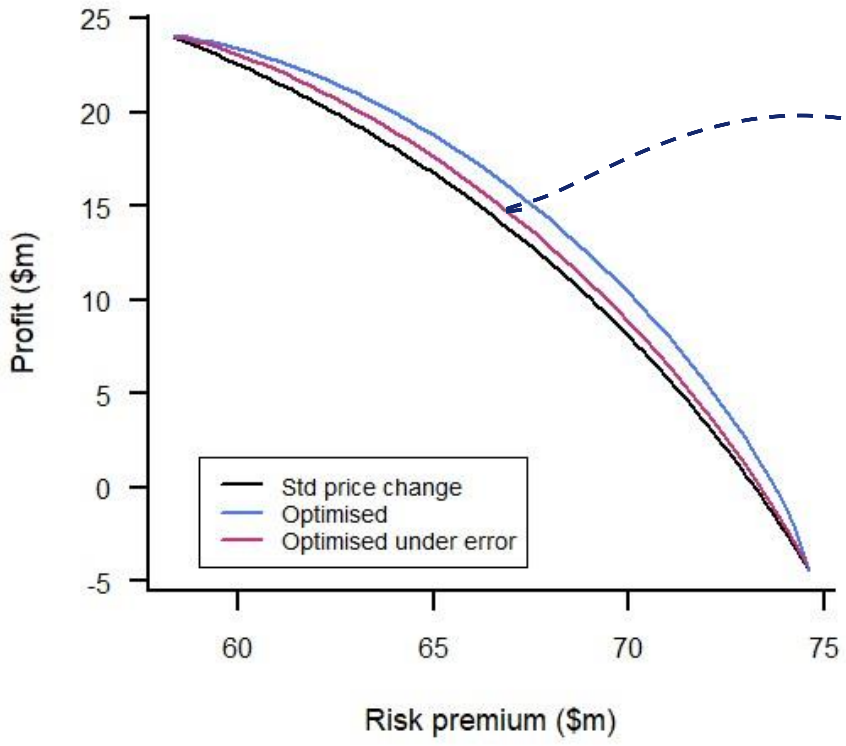

Relative performance of optimisation under error as T varies.

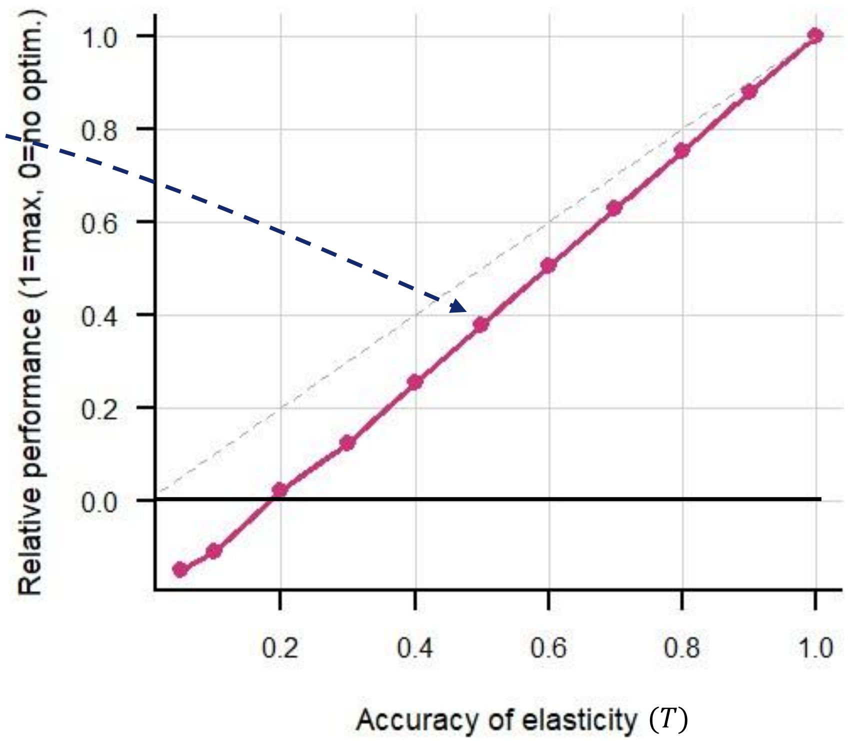

We observe, in our base setup:

- Significant (60%) degradation when $T = 0.5$. Importantly, while profit increase is comparable, customer volumes drop sharply.
- Worse than no-optimisation performance for low values of T. This is a new result.
- The slope of the red line in the second figure depends on the relative size of retention variability and elasticity variability.

Actuaries Institute.

Overconfidence in elasticity estimates (3)

Larger elasticity variability, smaller retention variability

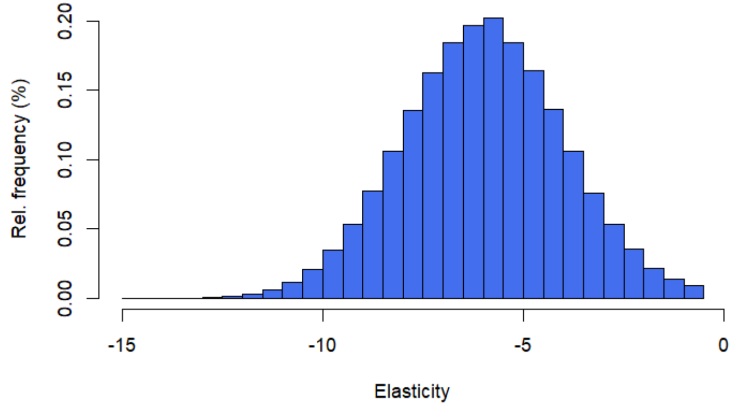

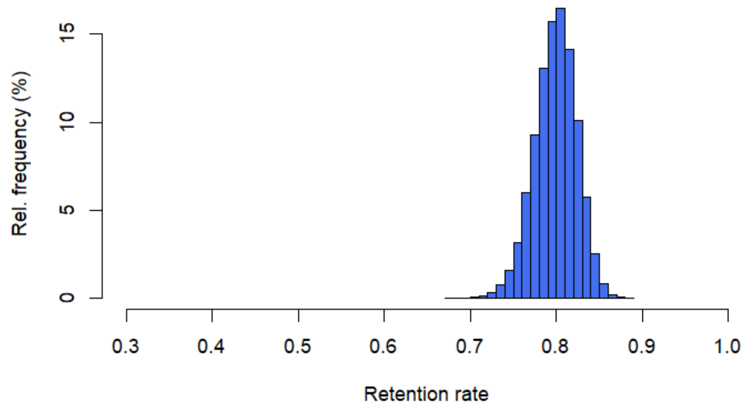

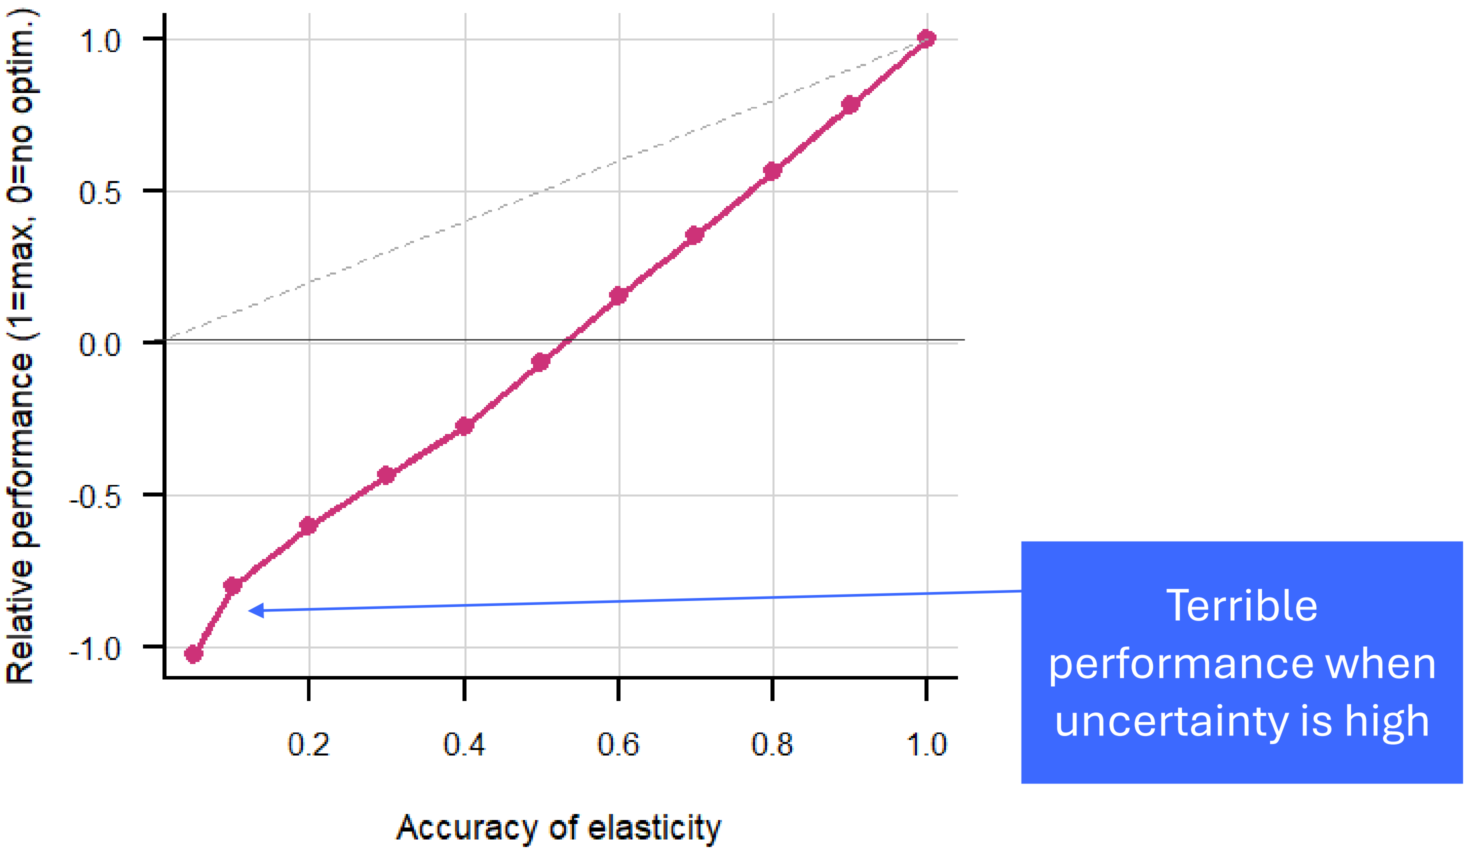

Actuaries Institute.

Overconfidence in elasticity estimates (4)

Smaller elasticity variability, larger retention variability

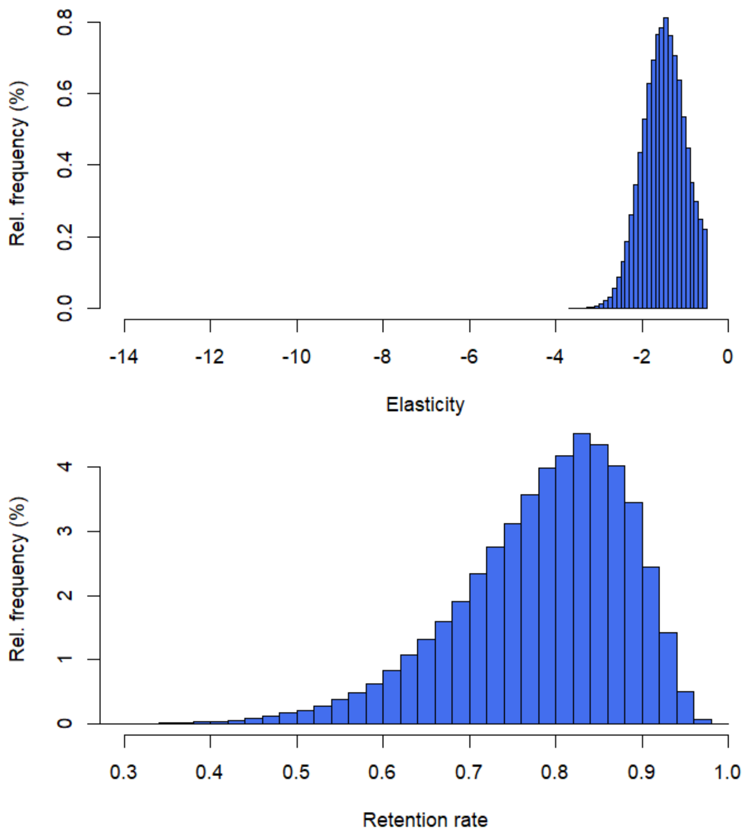

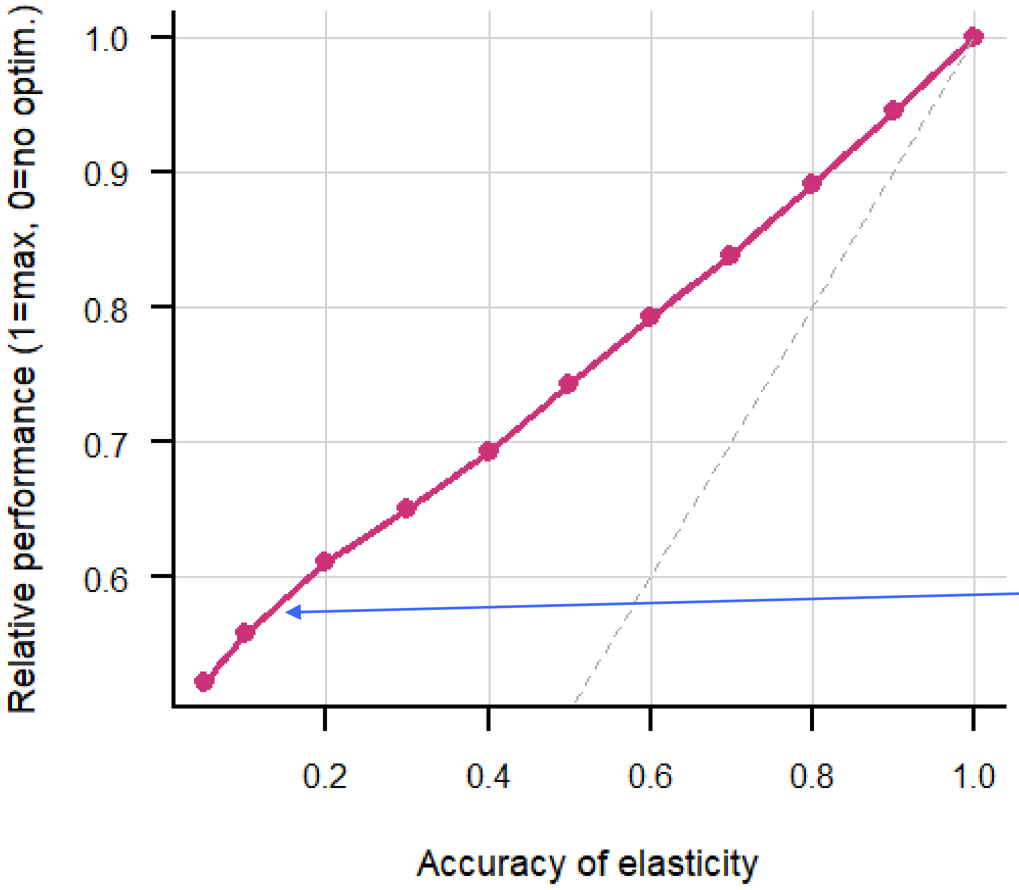

Decent performance, even when uncertainty is high

Variation in elasticity is hugely important. If true elasticity variation is low, it still may be possible to 'optimise on retention rates'. If true elasticity variation is high, poor estimation will lead to negative performance.

Actuaries Institute.

02 – Issues to consider when applying optimisation

A. Choice of constraints in optimisation

B. Drivers of optimised price changes

C. Overconfidence in elasticity estimates

D. Issues with model misspecification

E. Impact on the customer base over time

F. Challenges of modelling elasticity

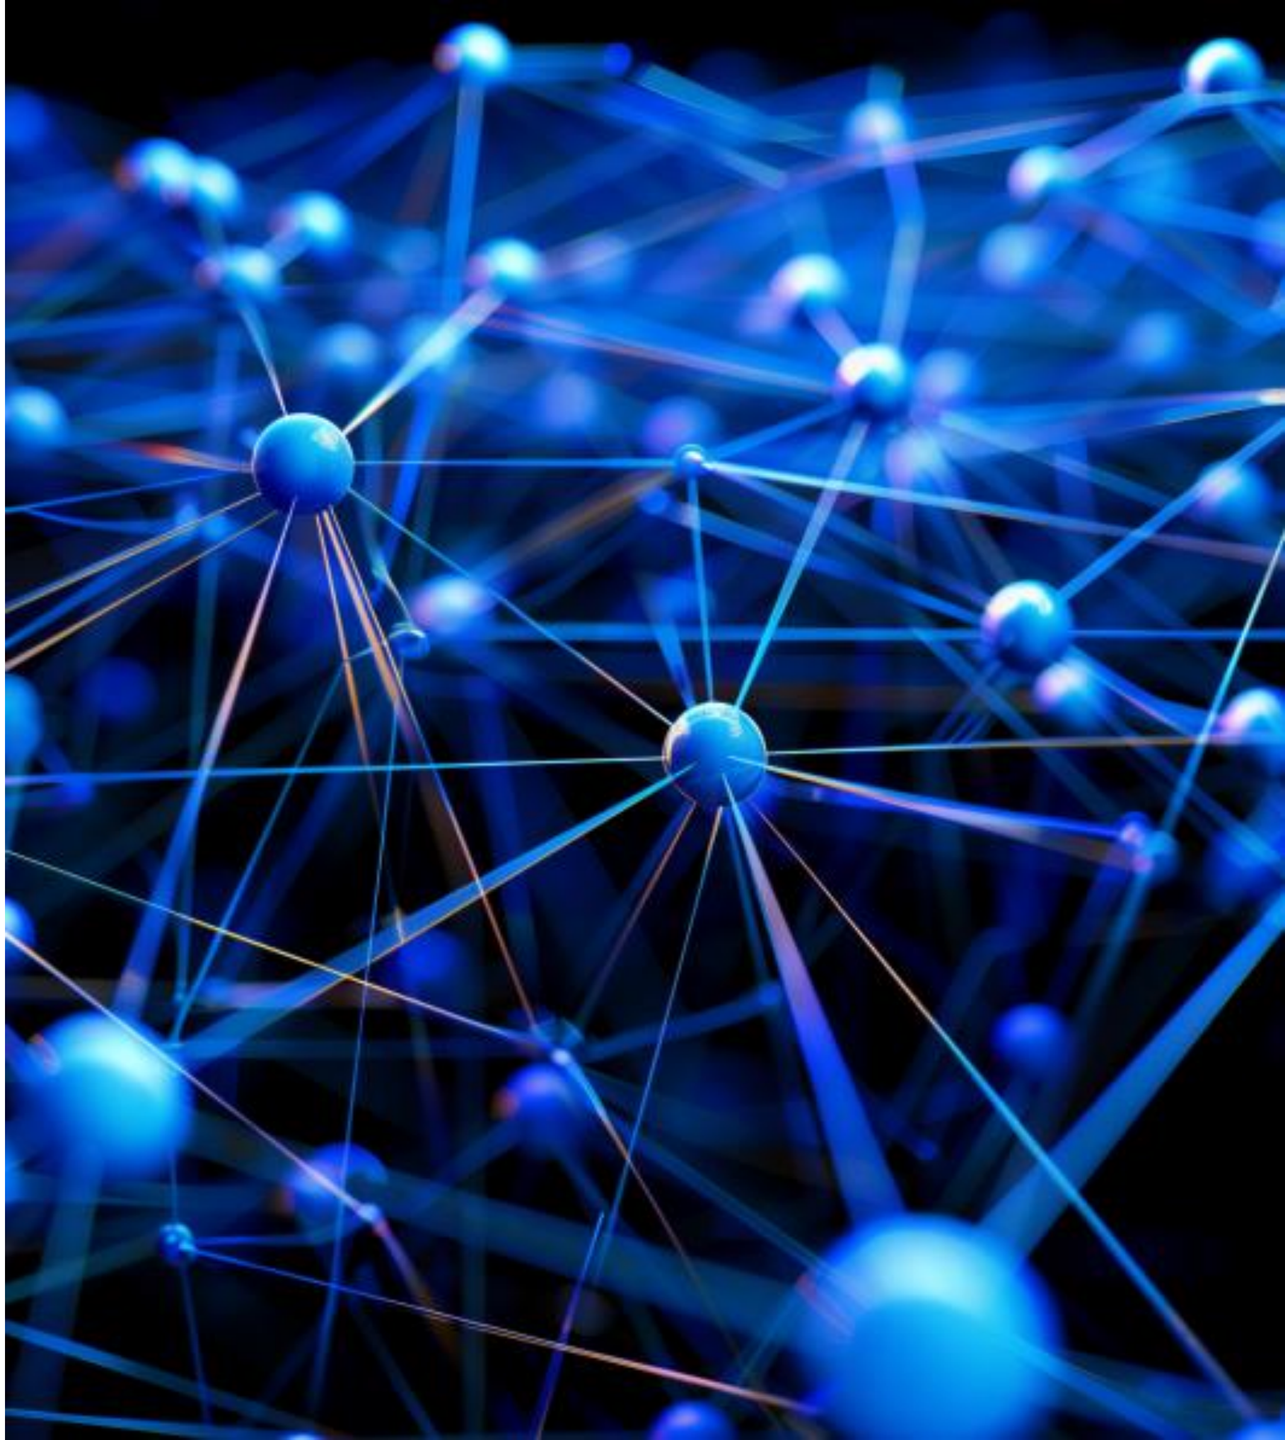

Actuaries Institute.

Issues with model misspecification

The preceding analysis assumes the model specification is correct, but elasticity parameter is subject to error.

We also consider misspecification of the model *structure*, even if estimates of elasticity are perfectly ordered.

In practice model validation checks can correct some of the issues with misspecification (e.g. adding interactions), but paucity of data may limit its extent.

We choose the best possible estimates of elasticity under our original model structure, where true elasticity follows alternative structures

Alternative 1 – Elasticity effect also depends on retention linear predictor

$$g(x_i) = g\big(z_i + e_i^* p_i (z_i + 1)\big)$$

Alternative 2 – Elasticity effect uses a probit link rather than logistic

$$g(x_i) = \Phi(z_i + e_i^* p_i)$$

Alternative 3 – Elasticity is on a 'per dollar' basis rather than a percentage basis

$$g(x_i) = g(z_i + e_i^* p_i P_i)$$

e^* tuned so that overall risk volume sensitivity (impact of a 5% price increase) is unchanged.

Actuaries Institute.

Issues with model misspecification (2)

Optimal versus misspecified price changes

Alternative 1 –
extra retention correlation

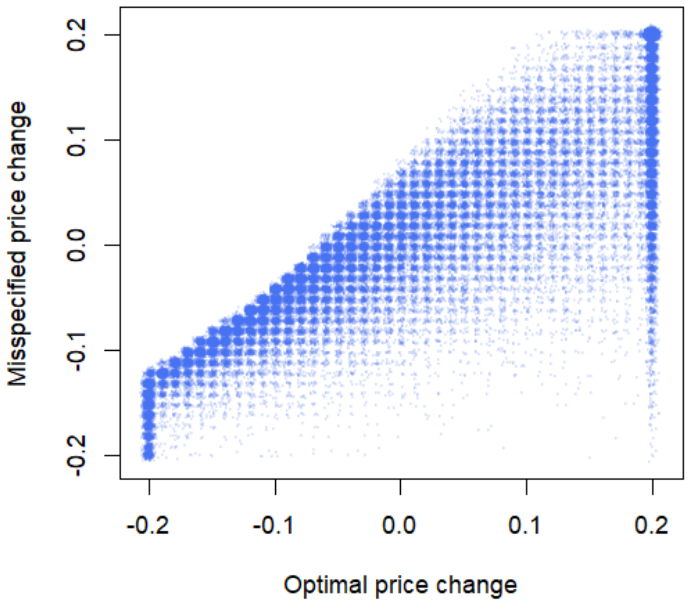

Alternative 2 –
Probit link

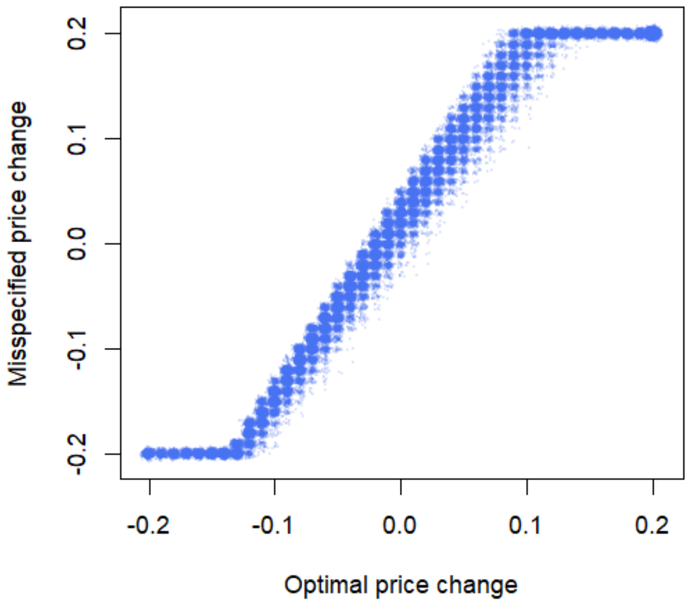

Alternative 3 –
Per dollar elasticity

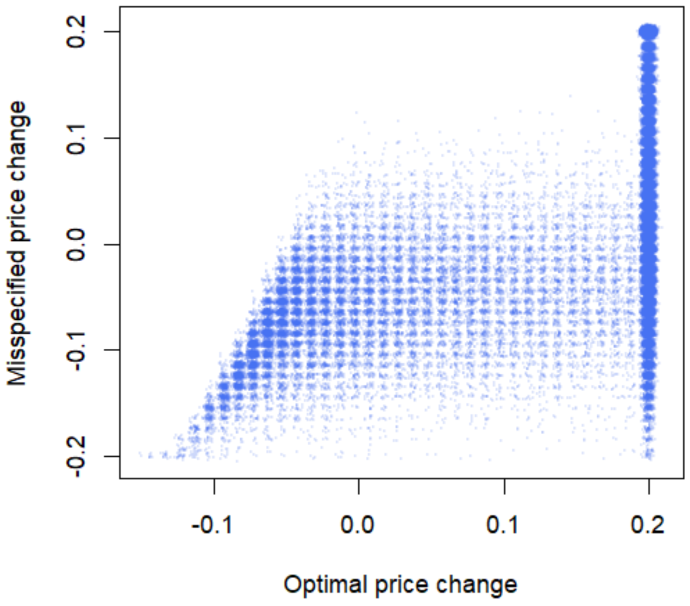

Misspecification of model structure can result in substantial changes in customer-level premiums

Actuaries Institute.

Issues with model misspecification (3)

Optimisation frontiers under optimal and misspecified price changes

Alternative 1 –
extra retention correlation

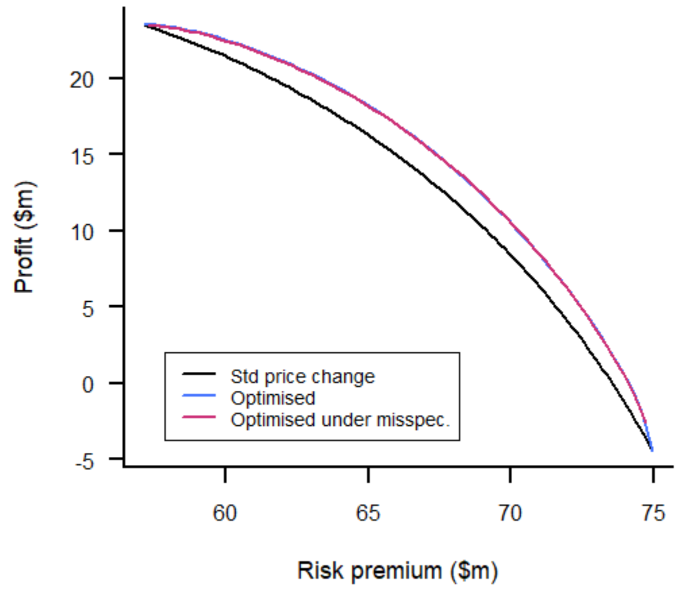

Alternative 2 –
Probit link

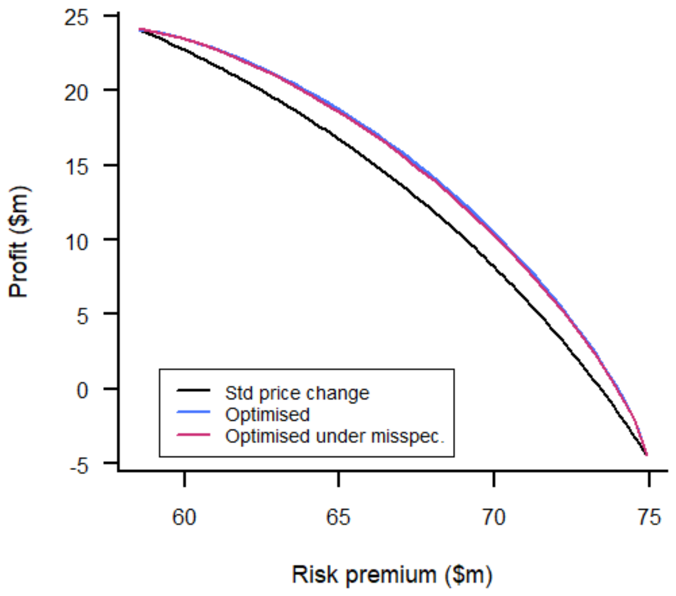

Alternative 3 –
Per dollar elasticity

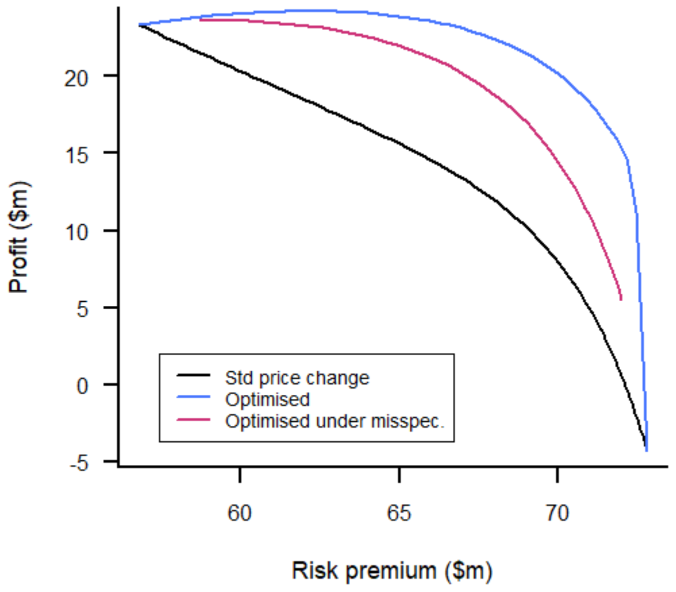

Actuaries
Institute.

02 – Issues to consider when applying optimisation

A. Choice of constraints in optimisation

B. Drivers of optimised price changes

C. Overconfidence in elasticity estimates

D. Issues with model misspecification

E. Impact on the customer base over time

F. Challenges of modelling elasticity

Actuaries Institute.

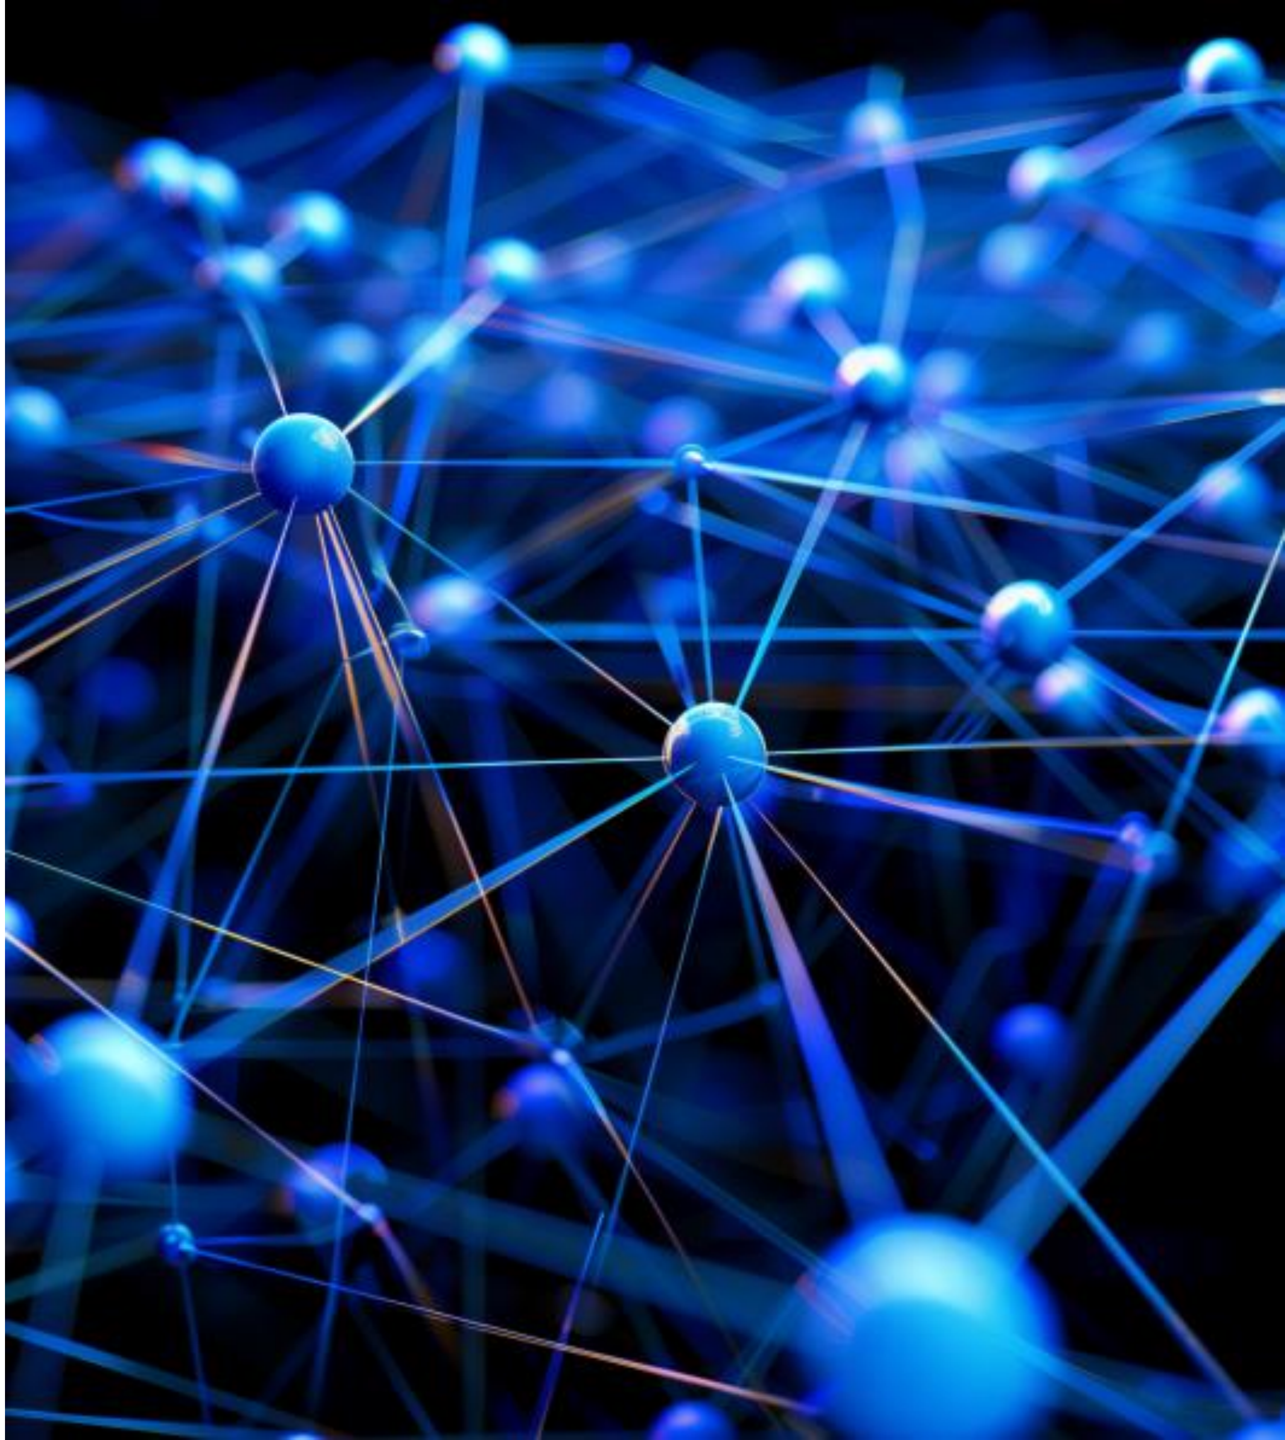

Impacts on customer base over time

We test a simplified multi-year setting, where we optimise over a single year, and adopt those prices for five years.

We compare the distribution of base retention and elasticity after 5 years, vs a uniform price change strategy.

After 5 years, optimisation results in a book with:

- Higher base churn rates (lower retention)
- Higher elasticity

Histograms showing change in customer base after five years of standard pricing and (simplified) optimisation. Green curve shows targeted profit margins at different level of elasticity and retention.

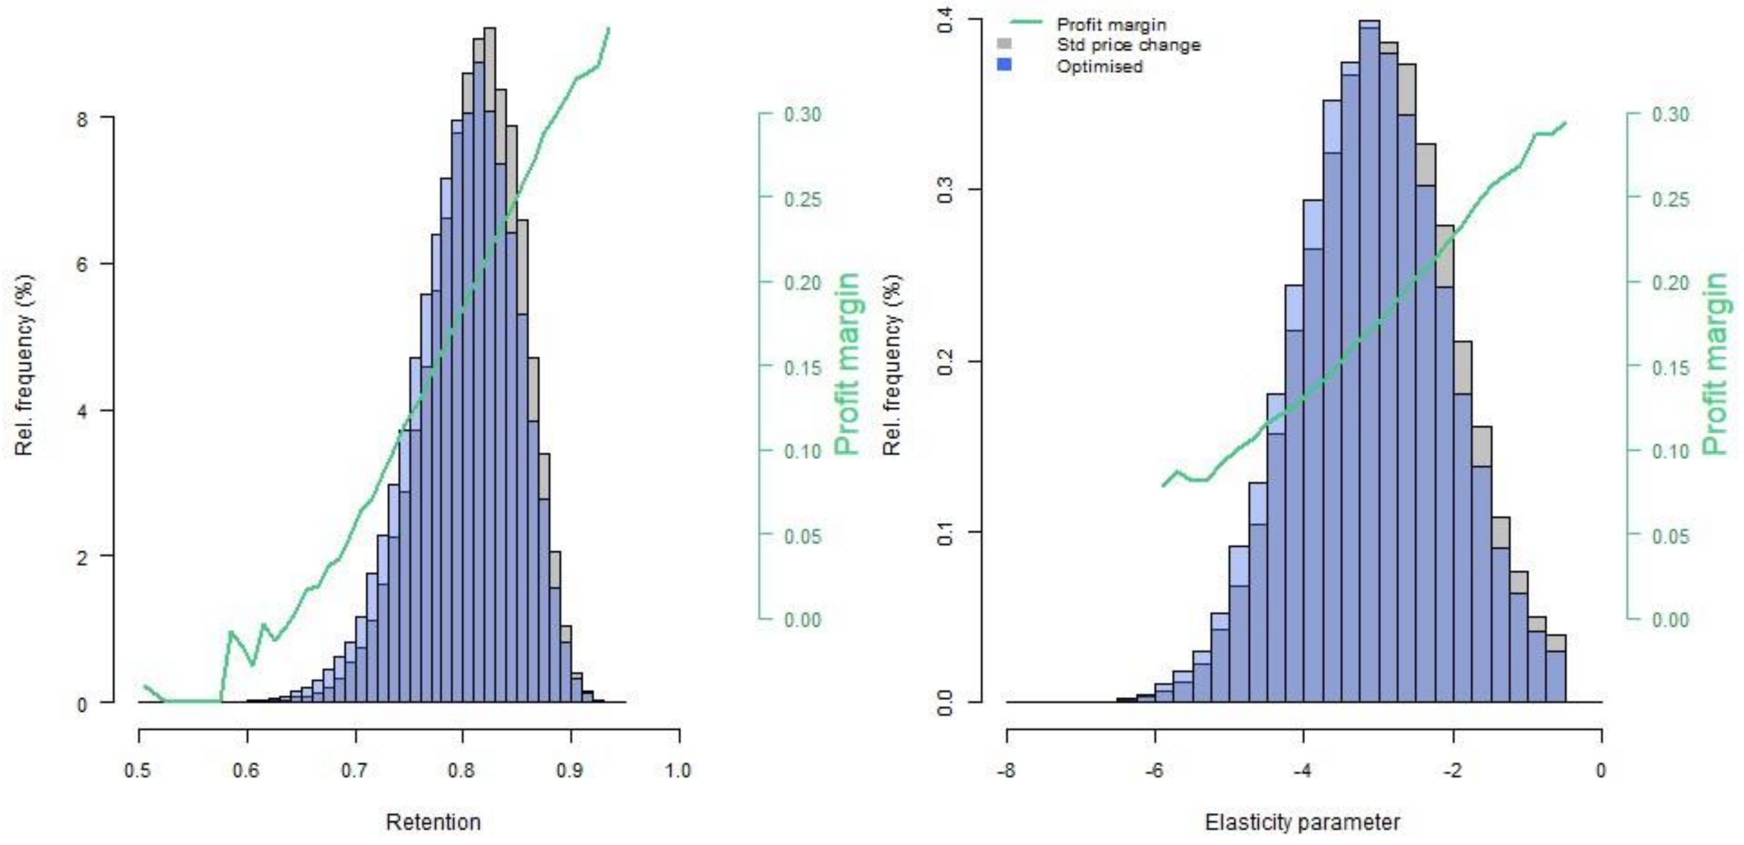

Actuaries Institute.

Impacts on customer base over time (2)

We compare profit from our simplified multi-year optimisation vs uniform pricing over time.

We consider scenarios where elasticity is:

- Known perfectly
- Misestimated (T = 0.5)
- Misestimated (T = 0.5) and true elasticity varies each year.

Optimisation gains decay towards zero over time.

Profit relatively compared to simple pricing strategies. Here 1 = similar profit to no-optimization targeting of the same volumes

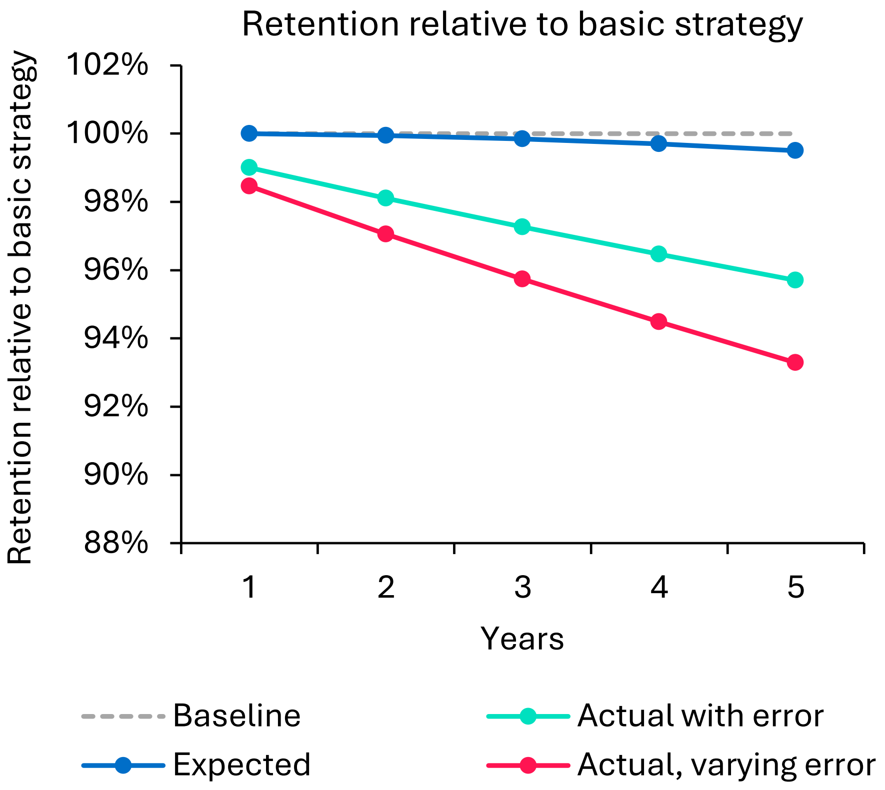

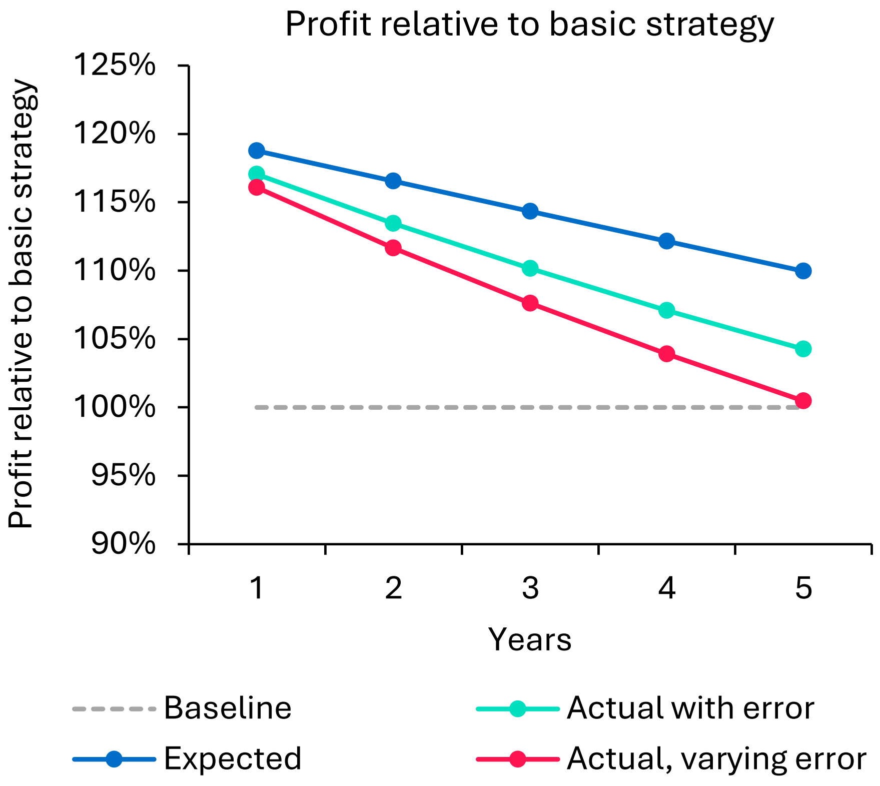

Actuaries Institute.

02 – Issues to consider when applying optimisation

A. Choice of constraints in optimisation

B. Drivers of optimised price changes

C. Overconfidence in elasticity estimates

D. Issues with model misspecification

E. Impact on the customer base over time

F. Challenges of modelling elasticity

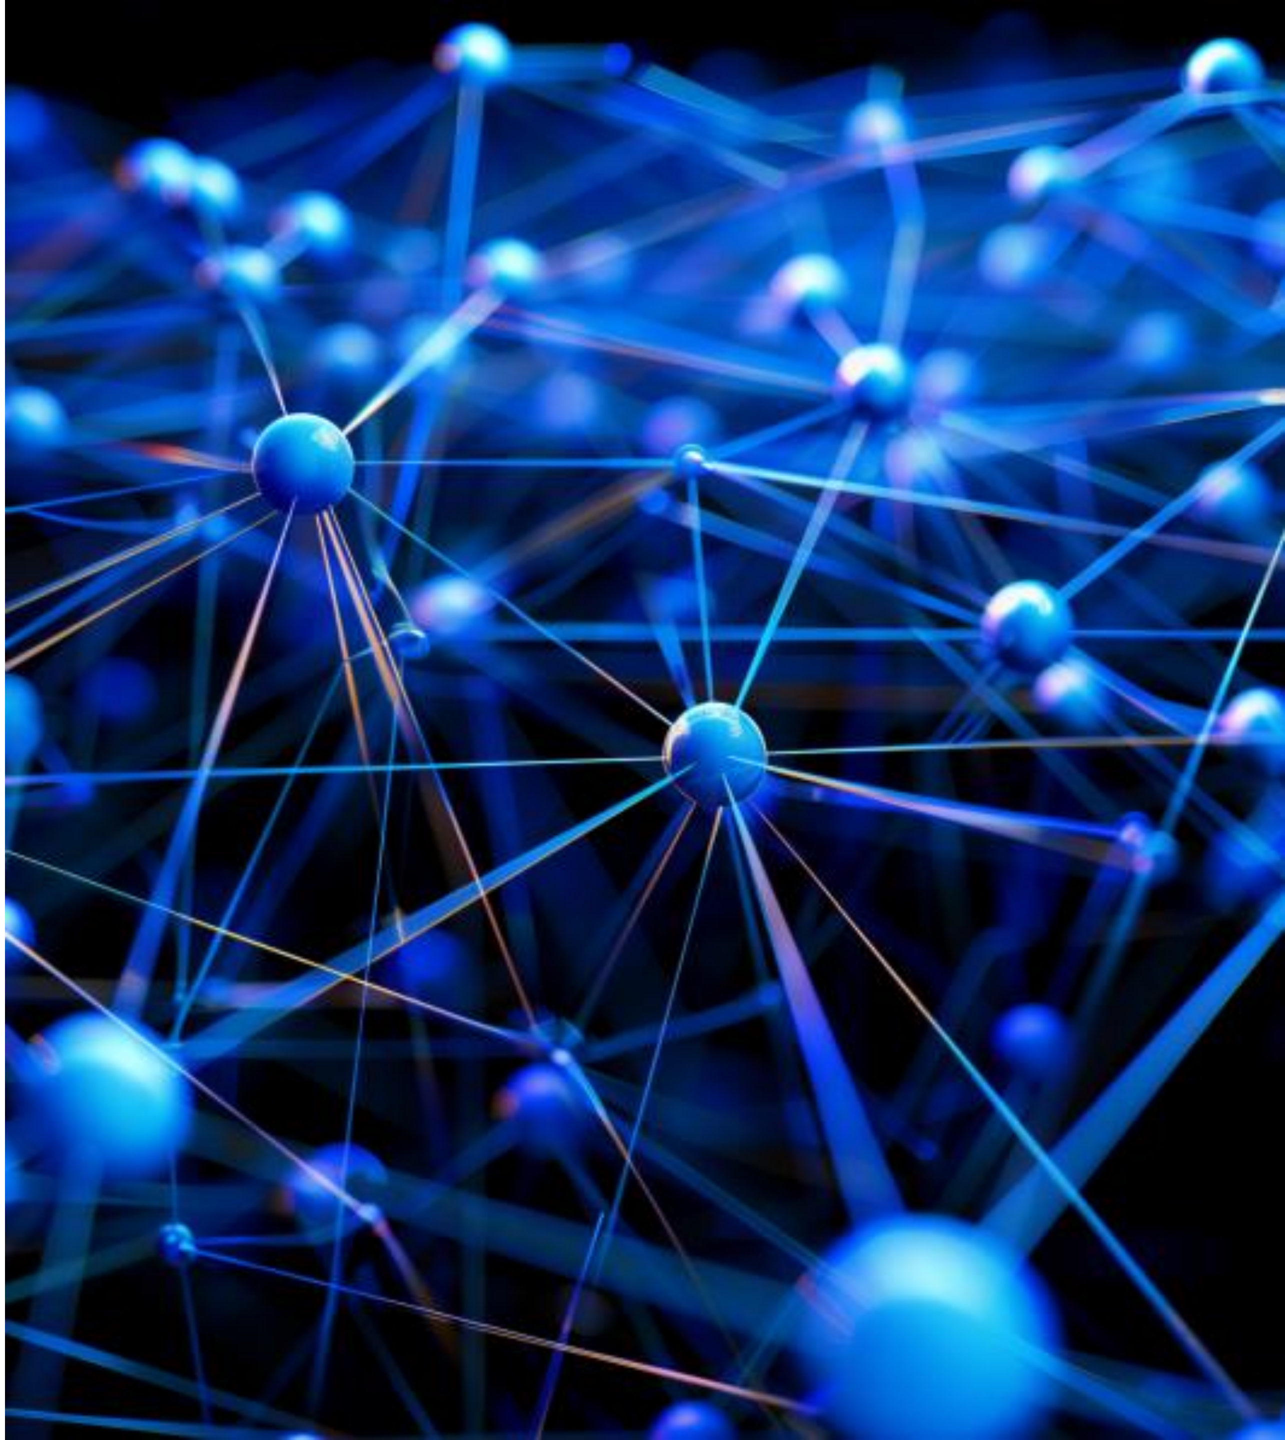

Actuaries Institute.

Challenges of modelling elasticity

Elasticity estimation is a causal modelling problem

- For optimisation, we want different elasticity estimates for different subgroups.
  - This is a heterogenous treatment effect modelling problem – harder again!
- This introduces the standard challenges associated with causal modelling, including:
  - Concerns around confounding variables
  - **A requirement for historical data that includes (ideally random) prices flexes** across the customer base.

Modelling approaches

- Machine learnings techniques exist to tackle these heterogenous treatment effect modelling problems (e.g. Athey et al. Causal Forests).
- However, tools to deploy these techniques are relatively immature.
- In practice, elasticity is often modelled as part of standard retention modelling, typically as a GLM.
  - Interactions with price change terms are used to capture elasticities for different cohorts.

Historical data with price flexes

Potential barriers to collecting price-flex data

- Technological costs and challenges
- Ethical or regulatory concerns
- Concerns around customer experience
- Lost profitability on policies that would otherwise have been optimised
- Distribution channel
- Portfolio size – is if feasible to collect enough?

Tempting alternatives

- Build elasticity models on data from historical price updates (e.g. risk model update).
  - Can be ok, but often limited systematic price flexing – biases in flex by rating factor
  - Fixed point in time.
- Build elasticity models on historically price optimised policies
  - Can work if well controlled, but need to be very careful about controlling for bias.
  - Can result in very poor outcomes if biases in price assignment are not adjusted for.
  - Less useful price exploration than at first glance.

Actuaries Institute.

Building elasticity models on optimised data (1)

We explore possible impacts of training elasticity models on optimised data in a simplified simulated setting.

We simulate model training data in two scenarios:

- With random price flexes between $\pm 20\%$
- With prices based on a one-year optimisation.

For each scenario we:

- Train a retention/elasticity GLM with all relevant main effects ($\boldsymbol{x_1}$ - $\boldsymbol{x_4}$) and a single price change term $\boldsymbol{e}$.
    - The model does not include price change interactions. This is a simplifying assumption!
- Run an optimisation, using each of these models to estimate demand / elasticity.

Actuaries Institute.

Building elasticity models on optimised data (2)

We show parameter estimates from the retention models trained on each dataset, and compare with (average) known true parameters from the simulated data.

- The model trained on price flex data accurately recreates the true parameters.
- The model trained on optimised data has a much lower estimate of elasticity.
  - This occurs due to correlations between other variables and price changes in training data.

True parameters and estimates from retention models

Parameter	Int.	x_1	x_2	x_3	x_4	e
True parameters	1.38 (average)	0.10	0.00	1.00	0.00	-3.00 (average)
Price flex data model	1.38	0.11	0.00	1.00	-0.04	-3.00
Optimised data model	1.44	0.04	0.00	0.75	0.11	-1.84

Training a model on historically optimised data can misestimate elasticity, if interactions are not perfectly specified.

Building elasticity models on optimised data (3)

We show results from running an optimisation using each retention model.

- The model trained on price flex data generates modest optimisation gains (+7%) vs a uniform price change.
- The model trained on optimised data is no better than a uniform price change.
  - It also results in a wider distribution of price changes, so could be viewed as worse.

Results from optimising with retention models

Scenario	Risk prem ($m)	Profit ($m)
Uniform price change	68.00	11.96
True parameters	68.00	14.25
Price flex data model	68.00	12.84
Optimised data model	68.00	11.96

If not properly specified, a model trained on optimised data can result in little (or negative!) benefit from optimisation.

Actuaries Institute.

Data volumes for elasticity estimation

Estimating interactions requires substantially more data than estimating main effects.

- Effect sizes are often smaller
- Estimating an effect for a sub-population means less data.

The size of historical price changes also heavily influences required data volumes.

- Halving price flexes = four times the data needed

Important considerations around price flex data:

- How much should be collected?
- How much price flex should be applied?

Actuaries Institute.

Data volumes for elasticity estimation

We simulate the impact of price flex magnitudes and data volumes on optimisation outcomes. We:

- Simulate data in our standard setup
- Randomly select price flexes between a given upper and lower limit, and simulate retention outcomes
- Down sample the dataset to the target size
- Build a retention model on this dataset
- Optimise premiums using this retention model.

We include interactions between price change and a binned version of true elasticity.

We repeat this process 50 times for each combination of price flex and sample size, and average across optimisation results.

Actuaries Institute.

Data volumes for elasticity estimation

Relationship between retention modelling data size and optimisation performance

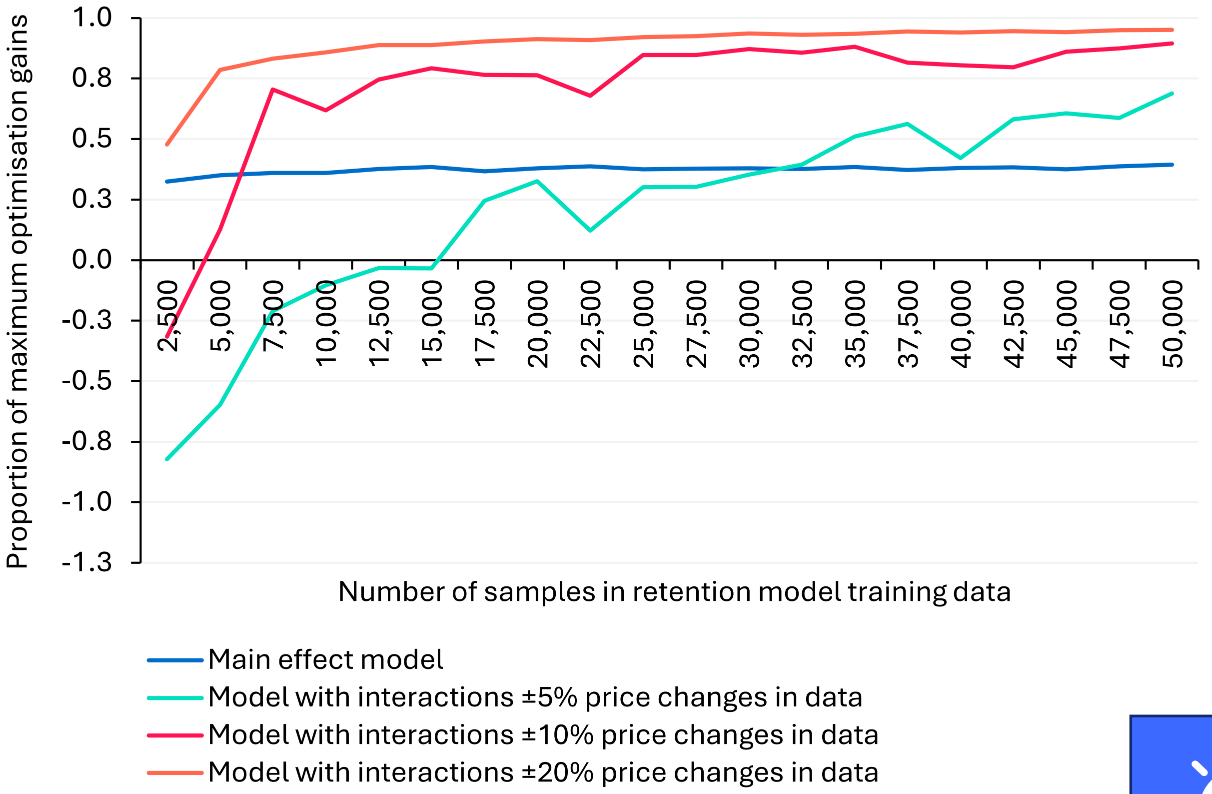

We observe:

- For small volumes of price flex data, including elasticity interactions can result in worse performance, due to misestimation.
- Approximately 5,000 – 10,000 samples of ±10% price flex data was needed to model simple interactions in this setup.
- Optimisation gains increase with sample size of price flexed data, with diminishing returns (roughly in relation to $\sqrt{n}$).
- Optimisation gains increase significantly with the magnitude of historical price flexes.

A substantial volume of randomised price flex data is needed to accurately model the elasticity interactions that power optimisation

Actuaries Institute.

Concluding thoughts

03

Actuaries Institute.

Key takeaways

There are several complexities in optimisation which mean that actual benefits may not be as high as estimated. Optimisation can provide value, but the following should be considered in the viability complex price optimisation:

1. It is easy to optimise on **inferior constraints** that lead to perverse, or at least unhelpful, portfolio outcomes
2. **Retention as well as elasticity** has a role to play in optimising premiums, and that the adopted link function as a significant impact on resulting premiums
3. **Overconfidence in elasticity estimation** leads to degradation in optimisation performance in a way that is not always obvious
4. Assumptions around **model structures have profound impacts** on results, when no ‘true’ model structure is known in advance
5. Optimisation leads to a **more elastic customer base over time**, which can carry longer-term implications
6. The ability to estimate customer elasticity (a core element of optimisation) relies on a good experimentation setup and testing volumes – which carries a significant cost

Actuaries Institute.

Thank you

Actuaries Institute
actuaries.asn.au