TESTING FOR MORE POSITIVE EXPECTATION DEPENDENCE WITH APPLICATION TO MODEL COMPARISON AND AUTOCALIBRATED MODELS

Michel Denuit (UCLouvain, Belgium)

Julien Trufin (ULB, Belgium)

Thomas Verdebout (ULB, Belgium)

OWARS Seminar, February 23, 2022

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

Outline

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

Positive expectation dependence

- Y and Z: 2 random variables.
- Y is **positively expectation dependent** on Z if

$$\begin{aligned} & \mathrm{E}[Y] \geq \mathrm{E}[Y|Z \leq z] \text{ for all } z \\ \Leftrightarrow\ & \mathrm{E}[Y|Z > z] \geq \mathrm{E}[Y] \text{ for all } z. \end{aligned}$$

Insurance pricing and model comparison

- Y: **response variable** (number of claims, claims severity).
- Z_1 and Z_2: **ranks of model predictions** under two competing insurance pricing tools.
- **Model comparison:**

  Model 1 outperforms model 2 if the response Y is more positively expectation dependent on Z_1 than on Z_2, that is, if

$$\begin{aligned} &\mathrm{E}[Y|Z_1 \leq z] \leq \mathrm{E}[Y|Z_2 \leq z] \text{ for all } z \\ \Leftrightarrow\ &\mathrm{E}[Y|Z_1 > z] \geq \mathrm{E}[Y|Z_2 > z] \text{ for all } z. \end{aligned}$$

Outline

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

Supervised learning

- A **response** Y and a **set of features** $\boldsymbol{X} = (X_1, \ldots, X_p)$.
- **Target:** conditional expectation $\mu(\boldsymbol{X}) = \mathrm{E}[Y|\boldsymbol{X}]$.
- $\mu(\boldsymbol{X})$: unknown and approximated by a **predictor** $\pi(\boldsymbol{X})$ **with a simpler structure**.
- **Assumption:** $\pi(\boldsymbol{X})$ is a continuous random variable.
  Notation:

$$F_\pi(t) = \mathrm{P}[\pi(\boldsymbol{X}) \leq t], \quad t \geq 0.$$

$$F_\pi^{-1}(\alpha) = \inf\{t \in \mathbb{R} | F_\pi(t) \geq \alpha\} \text{ for a probability level } \alpha.$$

Concentration curve

- **Predictor performances:** measured with the concentration curves.
- **Concentration curve** of $\mu(\boldsymbol{X})$ with respect to $\pi(\boldsymbol{X})$:

$$\mathrm{CC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X}); \alpha] = \frac{\mathsf{E}\big[\mu(\boldsymbol{X})\mathsf{I}[\pi(\boldsymbol{X}) \leq F_\pi^{-1}(\alpha)]\big]}{\mathsf{E}[\mu(\boldsymbol{X})]}.$$

$\Rightarrow$ CC assesses the dependence within the pair $(\mu(\boldsymbol{X}), \pi(\boldsymbol{X}))$.

- It turns out that

$$\mathrm{CC}\big[\mu(\boldsymbol{X}), \pi(\boldsymbol{X}); \alpha\big] = \mathrm{CC}[Y, \pi(\boldsymbol{X}); \alpha].$$

$\Rightarrow$ We can replace $\mu(\boldsymbol{X})$ with the response Y in CC.

- CC can be equivalently rewritten as

$$\mathrm{CC}[Y, \pi(\boldsymbol{X}); \alpha] \quad = \quad \frac{\mathsf{E}\big[Y \big| \pi(\boldsymbol{X}) \leq F_\pi^{-1}(\alpha)\big]}{\mathsf{E}[Y]} \times \alpha.$$

Model comparison

- $\pi_1(\boldsymbol{X})$ outperforms $\pi_2(\boldsymbol{X})$ if, and only if, $\forall\alpha \in (0,1)$

$$
\begin{aligned}
& \mathrm{CC}[Y, \pi_1(\boldsymbol{X}); \alpha] \leq \mathrm{CC}[Y, \pi_2(\boldsymbol{X}); \alpha] \\
\Leftrightarrow \quad & \mathsf{E}\big[Y \big| \pi_1(\boldsymbol{X}) \leq F_{\pi_1}^{-1}(\alpha)\big] \leq \mathsf{E}\big[Y \big| \pi_2(\boldsymbol{X}) \leq F_{\pi_2}^{-1}(\alpha)\big] \\
\Leftrightarrow \quad & \mathsf{E}\big[Y \big| F_{\pi_1}\big(\pi_1(\boldsymbol{X})\big) \leq \alpha\big] \leq \mathsf{E}\big[Y \big| F_{\pi_2}\big(\pi_2(\boldsymbol{X})\big) \leq \alpha\big] \\
\Leftrightarrow \quad & \mathsf{E}[Y | \Pi_1 \leq \alpha] \leq \mathsf{E}[Y | \Pi_2 \leq \alpha],
\end{aligned}
$$

where

$$
\Pi_1 = F_{\pi_1}\big(\pi_1(\boldsymbol{X})\big) \quad \text{and} \quad \Pi_2 = F_{\pi_2}\big(\pi_2(\boldsymbol{X})\big)
$$

are the corresponding ranks of $\pi_1(\boldsymbol{X})$ and $\pi_2(\boldsymbol{X})$.

$\Rightarrow$ **The ranking of the CC amounts to requiring that Y is more positively expectation dependent on Π_1 than on Π_2.**

Outline

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

Test $\mathcal{H}_0$ against $\mathcal{H}_1$

- We have **a random variable** Y **and two random variables** $\Pi_1 \sim Uni(0,1)$ **and** $\Pi_2 \sim Uni(0,1)$ possibly correlated between each other and with Y.
- **Observations:** n realizations $(Y_1, \Pi_{11}, \Pi_{21}), \ldots, (Y_n, \Pi_{1n}, \Pi_{2n})$.
- **Test:**

$$\mathcal{H}_0 : \mathrm{E}[Y|\Pi_1 \leq \alpha] \leq \mathrm{E}[Y|\Pi_2 \leq \alpha] \text{ for all } \alpha \in (0,1);$$
$$\mathcal{H}_1 : \mathrm{E}[Y|\Pi_1 \leq \alpha] > \mathrm{E}[Y|\Pi_2 \leq \alpha] \text{ for some } \alpha \in (0,1).$$

Π_1 and Π_2: identically distributed, so that $\mathcal{H}_0$ is equivalent to

$$\mathrm{E}[\mathsf{I}[\Pi_1 \leq \alpha]](\mathrm{E}[Y|\Pi_1 \leq \alpha] - \mathrm{E}[Y]) \leq \mathrm{E}[\mathsf{I}[\Pi_2 \leq \alpha]](\mathrm{E}[Y|\Pi_2 \leq \alpha] - \mathrm{E}[Y])$$

for all $\alpha \in (0,1)$, which in turn is equivalent to

$$\mathsf{C}\big[Y, \mathsf{I}[\Pi_1 \leq \alpha] - \mathsf{I}[\Pi_2 \leq \alpha]\big] \leq 0 \text{ for all } \alpha \in (0,1).$$

Test statistics T_n

- Let
$$D(\alpha) := \mathsf{C}\big[Y, \mathsf{I}[\Pi_1 \le \alpha] - \mathsf{I}[\Pi_2 \le \alpha]\big].$$
The most **natural estimator of $D(\alpha)$** is obtained by computing an **empirical covariance**, that is,
$$\widehat{D}(\alpha) := \frac{1}{n}\sum_{i=1}^{n}(Y_i - \bar{Y})(\mathsf{I}[\Pi_{1i} \le \alpha] - \mathsf{I}[\Pi_{2i} \le \alpha] - (\overline{\mathsf{I}[\Pi_1 \le \alpha]} - \overline{\mathsf{I}[\Pi_2 \le \alpha]}))$$
where $\bar{Y} := n^{-1}\sum_{i=1}^{n} Y_i$ and $\overline{\mathsf{I}[\Pi_k \le \alpha]} := n^{-1}\sum_{i=1}^{n}\mathsf{I}[\Pi_{ki} \le \alpha]$, $k = 1, 2$.
- We consider **a test that rejects** $\mathcal{H}_0$ at level $\beta \in (0,1)$ when
$$T_n := \sup_{\alpha\in(0,1)}\sqrt{n}\widehat{D}(\alpha) > \xi_\beta,$$
where the critical value ξ_β can be obtained by studying the limiting behavior of the empirical process $\sqrt{n}(\widehat{D}(\alpha) - D(\alpha))$ under $\mathcal{H}_0$.

Key result

- When $D(\alpha) = 0$ (at the boundary between $\mathcal{H}_0$ and $\mathcal{H}_1$), $\sqrt{n}\widehat{D}(\alpha)$ converges weakly to a **Gaussian process with mean zero and covariance function**

  $\Sigma(\alpha_1, \alpha_2) := \mathrm{E}\big[(Y - \mathrm{E}[Y])^2(\mathrm{I}[\Pi_1 \leq \alpha_1] - \mathrm{I}[\Pi_2 \leq \alpha_1])(\mathrm{I}[\Pi_1 \leq \alpha_2] - \mathrm{I}[\Pi_2 \leq \alpha_2])\big].$

- A Kolmogorov-Smirnov type **test can be obtained by rejecting $\mathcal{H}_0$ at the asymptotic level** $\beta \in (0, 1)$ when

  $$T_n = \sup_{\alpha \in (0,1)} \sqrt{n}\widehat{D}(\alpha) > c_\beta,$$

  where the critical value c_β can be obtained from the result above.

Outline

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

Data set

- **MTPL insurance portfolio** observed during one year.
- The portfolio includes **160 944 insurance policies**.
- For each policy i: the **numbers of claims** Y_i, the exposure-to-risk $e_i \leq 1$ (expressed in policy-year), and **eight features** $\boldsymbol{X}_i = (X_{i1}, \dots, X_{i8})$:
  - X_{i1} = AgePh: policyholder's age;
  - X_{i2} = AgeCar: age of the car;
  - X_{i3} = Fuel: fuel of the car, with two categories (gas or diesel);
  - X_{i4} = Split: splitting of the premium, with four categories (annually, semi-annually, quarterly or monthly);
  - X_{i5} = Cover: extent of the coverage, with three categories (from compulsory third-party liability cover to comprehensive);
  - X_{i6} = Gender: policyholder's gender, with two categories (female or male);
  - X_{i7} = Use: use of the car, with two categories (private or professional);
  - X_{i8} = PowerCat: the engine's power, with five categories.

Data set

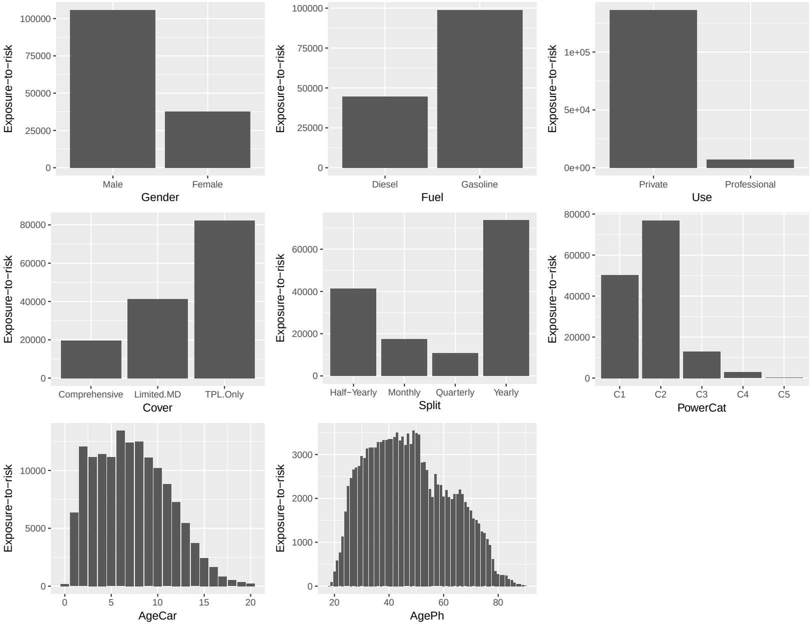
Exposure−to−risk
100000
75000
50000
25000
0
Male
Female
Gender
Exposure−to−risk
100000
75000
50000
25000
0
Diesel
Gasoline
Fuel
Exposure−to−risk
1e+05
5e+04
0e+00
Private
Professional
Use
Exposure−to−risk
80000
60000
40000
20000
0
Comprehensive
Limited.MD
TPL.Only
Cover
Exposure−to−risk
60000
40000
20000
0
Half−Yearly
Monthly
Quarterly
Yearly
Split
Exposure−to−risk
80000
60000
40000
20000
0
C1
C2
C3
C4
C5
PowerCat
Exposure−to−risk
10000
5000
0
0
5
10
15
20
AgeCar
Exposure−to−risk
3000
2000
1000
0
20
40
60
80
AgePh

Data set

Number of claims	Exposure-to-risk
0	126 499.7
1	15 160.4
2	1424.9
3	145.4
4	14.3
5	1.4
≥ 6	0

Table: Descriptive statistics for the number of claims.

Data set

- We partition the data set into a **training set** $\mathcal{D}$ and a **validation set** $\overline{\mathcal{D}}$.
- The **training set** $\mathcal{D}$ is composed of 80% of the observations taken at random from the entire data set.
- The **validation set** $\overline{\mathcal{D}}$ is made of the 20% remaining observations.

Models under consideration

- Y is assumed to be Poisson distributed with mean $e\mu(\boldsymbol{x})$.
  $\Rightarrow \mu(\boldsymbol{x}_i)$ represents the expected annual claim frequency for policyholder i.
- We aim to estimate the unknown function $\boldsymbol{x} \mapsto \mu(\boldsymbol{x})$.
- To that end, we first fit **2 GAMs** on $\mathcal{D}$:
  - $\pi^{\text{GAM1}}(\boldsymbol{x})$, with **only 2 features**: X_1 (AgePh) and X_2 (AgeCar);
  - $\pi^{\text{GAM2}}(\boldsymbol{x})$, using all **8 available features**.

  Notice that
  - The effects of AgePh and AgeCar: captured by splines.
  - No interaction terms.

Models under consideration

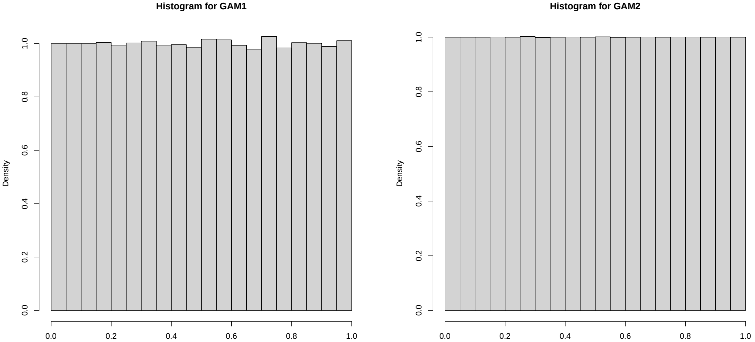

Figure: Histograms for Π^{GAM1} (left panel) and Π^{GAM2} (right panel) estimated from $\overline{\mathcal{D}}$.

Models under consideration

- Then, we fit **gradient boosting trees (GBT)**.
  - The **bagging fraction** $\gamma = 0.5$.
  - The **shrinkage** parameter $\tau = 0.01$.
  - The **size of the trees** is controlled by the interaction depth $\texttt{ID}$: $\texttt{ID} = 1, 2, 3, 4$.

  The training set $\mathcal{D}$ is divided into $\mathcal{D}_1$ (80%) and $\mathcal{D}_2$ (20%):
  - We train the GBT on $\mathcal{D}_1$.
  - We fine-tune the GBT (selection of the number of trees) on $\mathcal{D}_2$. We get $T = 1186, 1043, 633, 721$ for $\texttt{ID} = 1, 2, 3, 4$, respectively.

Models under consideration

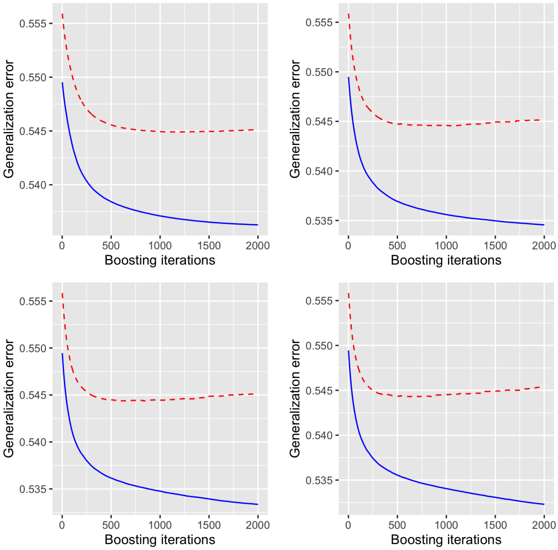

Figure: IS (blue) and OOS (red) of the gen. error for $\texttt{ID} = 1$ (top-left), $\texttt{ID} = 2$ (top-right), $\texttt{ID} = 3$ (bottom-left) and $\texttt{ID} = 4$ (bottom-right).

Models under consideration

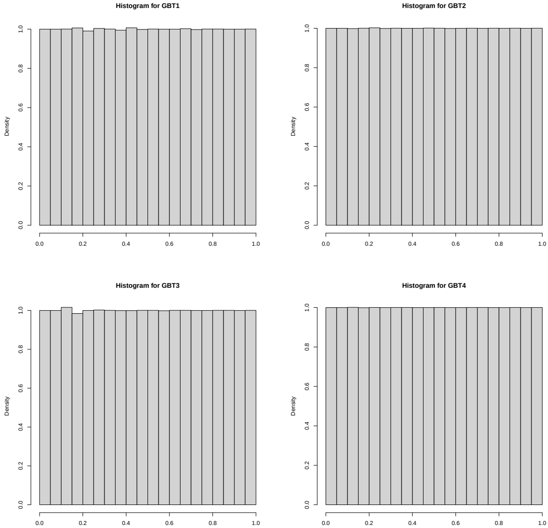

Figure: Distribution functions for Π^{GBT1} (top-left), Π^{GBT2} (top-right), Π^{GBT3} (bottom-left) and Π^{GBT4} (bottom-right) estimated on $\overline{\mathcal{D}}$.

Generalization errors

- **OOS estimates** (on $\overline{\mathcal{D}}$) of the **generalization errors** for the 6 models:

π^{GAM1}	$549.93 \cdot 10^{-3}$
π^{GAM2}	$548.05 \cdot 10^{-3}$
π^{GBT1}	$545.06 \cdot 10^{-3}$
π^{GBT2}	$544.54 \cdot 10^{-3}$
π^{GBT3}	$544.29 \cdot 10^{-3}$
π^{GBT4}	$544.30 \cdot 10^{-3}$

Testing procedure

- The **conditional expectation** $\mathrm{E}[Y|\Pi \leq \alpha]$ can be estimated on $\overline{\mathcal{D}}$ as

$$\widehat{\mathrm{E}}[Y|\Pi \leq \alpha] = \frac{\sum_{i \in \overline{\mathcal{D}}} y_i \mathrm{I}[\Pi(\boldsymbol{x}_i) \leq \alpha]}{\sum_{i \in \overline{\mathcal{D}}} \mathrm{I}[\Pi(\boldsymbol{x}_i) \leq \alpha]}.$$

Testing procedure

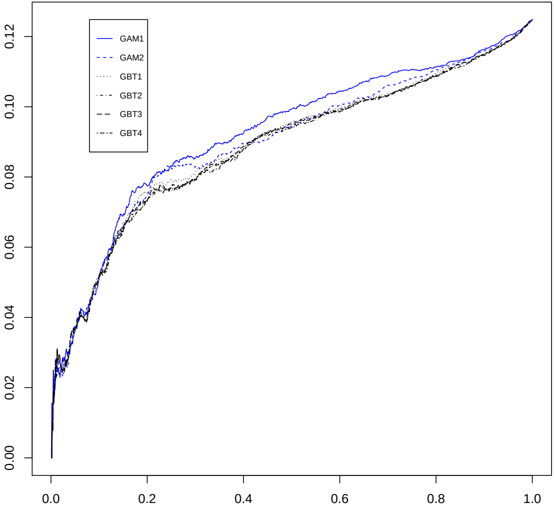

Testing procedure

		Π_2					
		Π^{GAM1}	Π^{GAM2}	Π^{GBT1}	Π^{GBT2}	Π^{GBT3}	Π^{GBT4}
Π_1	Π^{GAM1}	/	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	Π^{GAM2}	0.998	/	**0.049**	**0.008**	**0.022**	**0.010**
	Π^{GBT1}	1.000	0.710	/	0.420	0.256	0.232
	Π^{GBT2}	0.998	0.856	0.990	/	0.902	0.250
	Π^{GBT3}	1.000	0.616	1.000	0.792	/	0.230
	Π^{GBT4}	0.998	0.806	1.000	0.910	0.964	/

Table: Values of $\widehat{p}$. $\mathcal{H}_0 : \mathrm{E}[Y|\Pi_1 \leq \alpha] \leq \mathrm{E}[Y|\Pi_2 \leq \alpha]$ for all $\alpha \in (0,1)$ is rejected at the level 0.05 when $\widehat{p} < 0.05$ (cases printed in bold).

Conclusion

- $\mathcal{H}_0 : \mathrm{E}[Y|\Pi_1 \leq \alpha] \leq \mathrm{E}[Y|\Pi_2 \leq \alpha]$ for all $\alpha \in (0,1)$ is **always rejected for $\Pi_1 = \Pi^{\mathbf{GAM1}}$ whatever Π_2**.
- For $\Pi_1 = \Pi^{\mathrm{GAM2}}$, the same observation holds at the level 0.05 except for $\Pi_2 = \Pi^{\mathrm{GAM1}}$. This shows that **GBT outperforms GAMs** on this data set.
- The testing procedure **does not identify one GBT dominating the others**. This confirms the similar performances of all GBTs on $\overline{\mathcal{D}}$.

Outline

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

Autocalibrated predictors and Bregman dominance

- A predictor π is **autocalibrated** if $\pi(\boldsymbol{X}) = \mathrm{E}[Y|\pi(\boldsymbol{X})]$.
- **Bregman dominance**:
  π_1 outperforms π_2 in terms of Bregman dominance if the inequality $\mathrm{E}[L(Y, \pi_1)] \leq \mathrm{E}[L(Y, \pi_2)]$ holds true for every Bregman loss function L.
- **Bregman dominance and autocalibrated predictors**:
  π_1 outperforms π_2 in terms of Bregman dominance if, and only if,

$$\pi_2(\boldsymbol{X}) \preceq_{\mathrm{cx}} \pi_1(\boldsymbol{X}).$$

Autocalibrated predictors and Bregman dominance

- For **autocalibrated** predictors:

$$
\begin{aligned}
& \pi_2(\boldsymbol{X}) \preceq_{\mathsf{cx}} \pi_1(\boldsymbol{X}) \\
\Leftrightarrow \quad & \mathsf{LC}[\pi_1(\boldsymbol{X}); \alpha] \leq \mathsf{LC}[\pi_2(\boldsymbol{X}); \alpha] \text{ for all } \alpha \in (0,1) \\
\Leftrightarrow \quad & \mathsf{CC}[\mu(\boldsymbol{X}), \pi_1(\boldsymbol{X}); \alpha] \leq \mathsf{CC}[\mu(\boldsymbol{X}), \pi_2(\boldsymbol{X}); \alpha] \text{ for all } \alpha \in (0,1)
\end{aligned}
$$

  since Lorenz and concentration curves coincide for autocalibrated predictors.

- **In conclusion:**
  For **autocalibrated** predictors, the **testing procedure** developed in the present study is a **tool to test for Bregman dominance**.

Autocalibrated predictors and Tweedie dominance

- **Actuarial studies:** often assume distribution belonging to the **Tweedie subclass of the Exponential Dispersion family**
- Variance function of the form $V(\mu) = \mu^{\xi}$ for some power parameter ξ.
- $\xi \geq 1$: relevant cases for applications in insurance:
  - $\xi = 1$: **Poisson** distribution;
  - $1 < \xi < 2$: **Compound Poisson** sums with Gamma-distributed terms;
  - $\xi = 2$: **Gamma** distribution;
  - $\xi = 3$: **Inverse Gaussian** distribution.

Autocalibrated predictors and Tweedie dominance

- **Tweedie deviance** essentially reduces to

$$D(\xi,\pi)=\begin{cases}\mathrm{E}\left[\pi(\boldsymbol{X})-Y\ln\pi(\boldsymbol{X})\right] & \text{for } \xi=1\\ \mathrm{E}\left[\ln\pi(\boldsymbol{X})+\frac{Y}{\pi(\boldsymbol{X})}\right] & \text{for } \xi=2\\ \mathrm{E}\left[\frac{\pi(\boldsymbol{X})^{2-\xi}}{2-\xi}-\frac{Y\pi(\boldsymbol{X})^{1-\xi}}{1-\xi}\right] & \text{for } \xi>1 \text{ and } \xi\neq 2\end{cases}$$

- **Tweedie dominance**: π_1 outperforms π_2 in terms of Tweedie dominance if the inequality $D(\xi,\pi_1)\leq D(\xi,\pi_2)$ holds true for every power parameter $\xi\geq 1$.

Autocalibrated predictors and Tweedie dominance

- For two autocalibrated predictors π_1 and π_2, π_1 outperforms π_2 in terms of Tweedie dominance if, and only if, the inequality
$$\mathsf{E}\left[\psi_\xi(\pi_2(\boldsymbol{X}))\right] \leq \mathsf{E}\left[\psi_\xi(\pi_1(\boldsymbol{X}))\right]$$
holds true for every power parameter $\xi \geq 1$, where
$$\psi_\xi(\pi) = \begin{cases} \pi \ln \pi \text{ for } \xi = 1 \\ -\ln \pi \text{ for } \xi = 2 \\ \frac{\pi^{2-\xi}}{\xi-2} \text{ for } \xi > 1 \text{ and } \xi \neq 2. \end{cases}$$

Autocalibrated predictors and Tweedie dominance

- Let $\mathrm{L}_\pi(\cdot)$ denote the Laplace transform of π, that is, $\mathrm{L}_\pi(s) = \mathrm{E}[\exp(-s\pi)]$. For two autocalibrated predictors π_1 and π_2, if the inequality $\mathrm{L}_{\pi_2}(s) \leq \mathrm{L}_{\pi_1}(s)$ holds true for all $s \geq 0$, then π_1 outperforms π_2 in terms of Tweedie dominance.
- **Laplace order is a sufficient condition for Tweedie dominance**.

ICC and ABC metrics

- Denuit et al. (2019) suggest basing the comparison both on the **integral of the concentration curves (ICC)** and the **area between the two curves (ABC)**:

$$\mathrm{ICC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X})] = \int_0^1 \mathrm{CC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X}); \alpha] \mathrm{d}\alpha$$

  and

$$\mathrm{ABC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X})] = \int_0^1 \Big(\mathrm{CC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X}); \alpha] - \mathrm{LC}[\pi(\boldsymbol{X}); \alpha] \Big) \mathrm{d}\alpha.$$

- A better model has smaller ICC and ABC.

Gini coefficient

- Dominant practice uses **Gini coefficients** to compare predictors:

$$\mathrm{Gini}[Z] = \mathsf{E}\big[|Z_1 - Z_2|\big] = \mathsf{E}\big[\max\{Z_1, Z_2\}\big] - \mathsf{E}\big[\min\{Z_1, Z_2\}\big]$$

  where Z_1 and Z_2 are independent and distributed as Z.
- If Z is continuous then it can be shown that

$$\mathrm{Gini}[Z] = 4\mathsf{C}\big[Z, F_Z(Z)\big].$$

ICC, ABC and Gini coefficient for autocalibrated models

- If π is **autocalibrated**, then we have

$$\mathrm{ICC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X})] = \frac{1}{2} - \frac{\mathrm{Gini}[\pi(\boldsymbol{X})]}{\mathsf{E}[\pi(\boldsymbol{X})]}$$

  and $\mathrm{ABC}[\mu(\boldsymbol{X}), \pi(\boldsymbol{X})] = 0$.

- For autocalibrated models, π_1 outperforms π_2 when

$$\mathrm{ICC}[\mu(\boldsymbol{X}), \pi_1(\boldsymbol{X})] \leq \mathrm{ICC}[\mu(\boldsymbol{X}), \pi_2(\boldsymbol{X})]$$

  $\Leftrightarrow$

$$\mathrm{Gini}\big[\pi_1(\boldsymbol{X}), \Pi_1\big] \geq \mathrm{Gini}\big[\pi_2(\boldsymbol{X}), \Pi_2\big].$$

  This **justifies the dominant practice** provided autocalibration has been implemented.

Outline

Introduction and motivation

Expectation dependence in supervised learning

Testing for more positive expectation dependence

Case study

Autocalibrated predictors

References

References

- Denuit, M., Charpentier, A., Trufin, J. (2021a). **Autocalibration and Tweedie-dominance for insurance pricing with machine learning.** *Insurance: Mathematics and Economics*, 101 (Part B), 485-497.
- Denuit, M., Sznajder, D., Trufin, J. (2019). **Model selection based on Lorenz and concentration curves, Gini indices and convex order.** *Insurance: Mathematics and Economics* 89, 128-139.
- Denuit, M., Trufin, J., Verdebout, T. (2021b). **Testing for more positive expectation dependence with application to model comparison**. *Insurance: Mathematics and Economics* 101 (Part B), 163-172.