DRIFTS, SHIFTS AND INSTABILITIES IN SCORECARD MODEL RISK

Alan Forrest

Business Associate

Credit Research Centre, University of Edinburgh

Andrija Djurovic

External Consultant

Deloitte Sweden

CRC XIX

Edinburgh, August 2025

Disclaimer

This presentation is the authors’ own opinion and is not necessarily the view of the authors’ employers.

Any data or illustrations used in this presentation are not based on confidential business information or on personal data.

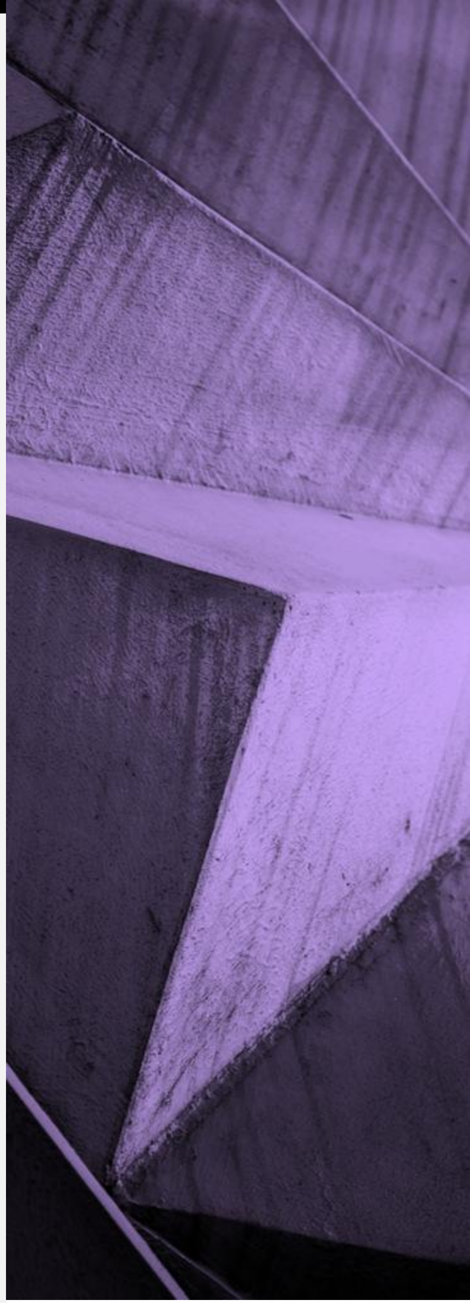

DRIFTS AND SHIFTS

.

- Data Drift is not always a bad thing - it’s just misunderstood :
  - And it’s poorly measured by blunt metrics like PSI.

- Data Drift can be quantified and split into benign components and damaging components:
  - A validation tool;
  - Practical model risk monitoring and early warning.

- Data Drift is channelled into useful Data Shift and Model Shift.

THREE KINDS OF DATA DRIFT

- Population Drift / Simple Covariate Shift.
- Label Shift / Prior Probability Shift / Change of Statistical Law.
- Concept Drift / Change of Model Structure.

Seen generally as a problem to be measured and managed:

> Measured by … PSI or other distribution comparisons on inputs and on output score…
> Managed by … performance monitoring, model rebuild process…

Can we take control and pre-empt damaging Data Drift, by quantifying it and using it to our advantage?

DETECTING DRIFT

An example dataset and model – imbalanced and poorly specified

Logistic regression scorecard model based on the following data

Base Data

		A2		
	Default / n	1	2	3
A1	1	50 / 1000	50 / 100	20 / 100
	2	5 / 100	5 / 1000	2 / 1000

$$d/n \sim A1 + A2$$

Scorecard = log odds bad

Int	A1=1	A2=1	A2=2
-3.731	2.250	-1.364	1.310

- Imbalanced
- High collinearity of factors
- Strong interaction, not specified in model
- Low default numbers.

DETECTING DRIFT

PSI will not detect reliably when Data Drift requires a model change

Base

	d/n	A2 = 1	A2 = 2	A2 = 3
A1	1	50/1000	50/100	20/100
	2	5/100	5/1000	2/1000

Scorecard

Int	A1=1	A2=1	A2=2
-3.731	2.250	-1.364	1.310

PSI*1000

A1	269.6
A2	688.3
A1 x A2	751.2

Shift 1

	d/n	A2 = 1	A2 = 2	A2 = 3
A1	1	47/992	48/99	25/104
	2	7/112	6/1001	12/6010

Scorecard

Int	A1=1	A2=1	A2=2
-3.742	2.271	-1.386	1.294

PSI*1000

A1	0.083
A2	0.713
A1 x A2	4.663

Shift2

	d/n	A2 = 1	A2 = 2	A2 = 3
A1	1	69/1018	40/98	5/89
	2	47/142	9/1004	5/1003

Scorecard

Int	A1=1	A2=1	A2=2
-3.255	0.233*	0.873	0.370

Can we do better?

SHIFTING DATA AND MODELS

Data and Models are both lists of count values – points in a high dimensional space

Scorecard logistic regression - with its chosen design - is equivalent to a log-linear model on the contingency table – with its corresponding design (Z = default outcome).

$$d/n \sim A1 + A2 \qquad \text{count} \sim A1 * A2 + Z * (A1 + A2)$$

Our implementation prefers the log-linear design because it treats population marginals and model outcomes equally. This is important to allow for all kinds Data Shift.

x Data Observed Counts

		A2		
	d/n	1	2	3
A1	1	50 / 1000	50 / 100	20 / 100
	2	5 / 100	5 / 1000	2 / 1000

y Model Expected Counts

		A2		
	d	1	2	3
A1	1	54.94	45.70	19.34
	2	0.06	9.28	2.66

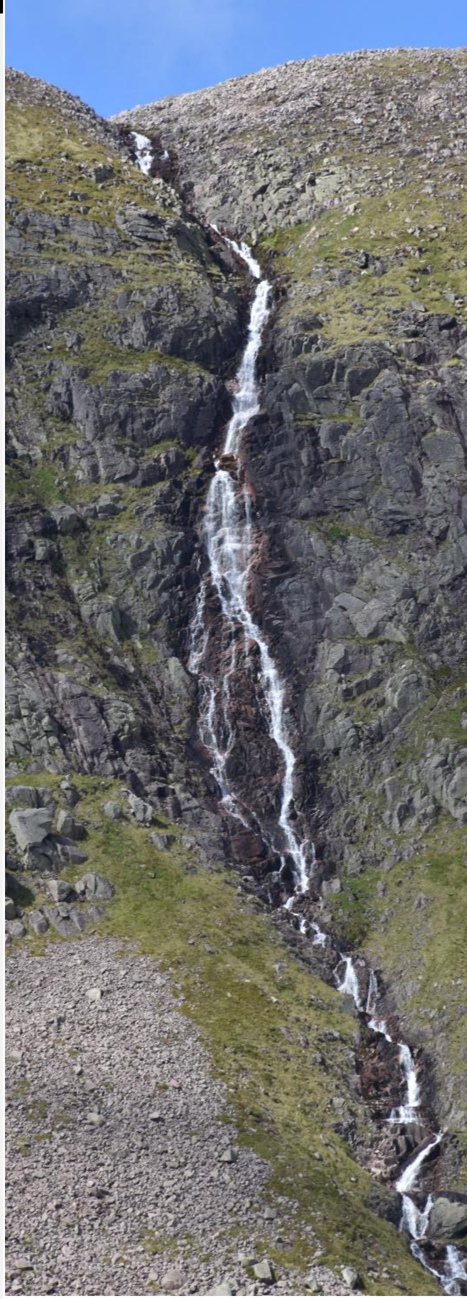

SHIFTING DATA AND MODELS

The data point can be shifted – what does the shifted (i.e. refitted) model look like?

Data = Observed

	A2		
d/n	1	2	3
A1 1	50 / 1000	50 / 100	20 / 100
2	5 / 100	5 / 1000	2 / 1000

$x \rightarrow x + \delta x$

Shifted Data

	A2		
d/n	1	2	3
A1 1	60 / 1010	45 / 99	13 / 95
2	26 / 121	7 / 1002	4 / 1002

Model = Expected

	A2		
d	1	2	3
A1 1	54.94	45.70	19.34
2	0.06	9.28	2.66

$y \rightarrow y + \delta y$

Shifted Model

	A2		
d/n	1	2	3
A1 1	??	??	??
2	??	??	??

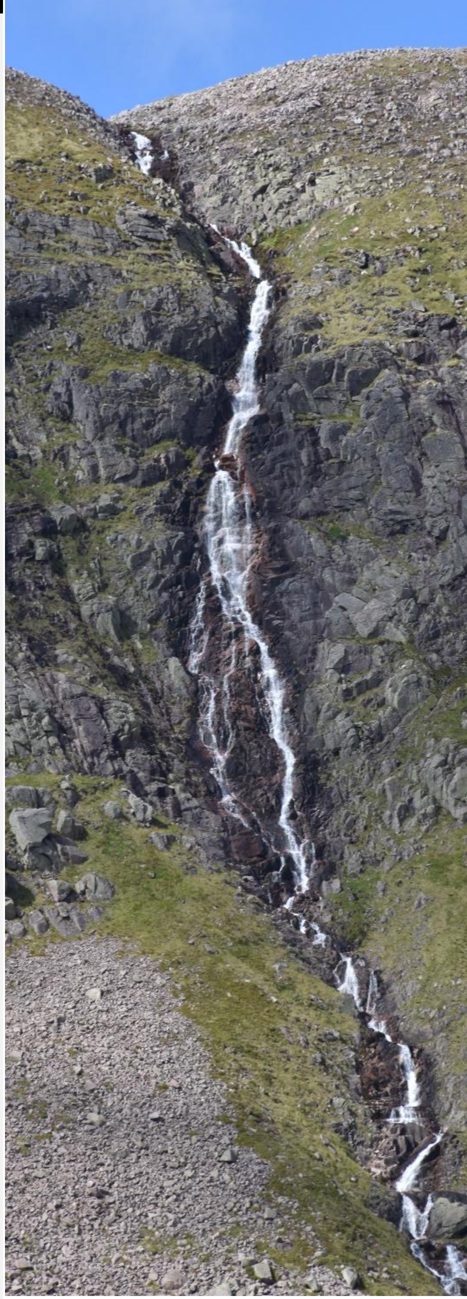

SHIFTING DATA AND MODELS

Let's also track what happens to the shifted model's parameters and its performance

Scorecard

Int	A1=2	A2=2	A2=3
-2.8451	-4.4994	2.6736	1.4171

b → $b + \delta b$

Scorecard for shifted model

Int	A1=2	A2=2	A2=3
??	??	??	??

Your chosen performance metric W on the base model → Performance of the shifted model $W + \delta W$

Can we go straight from δx to δy , δb , δW ?

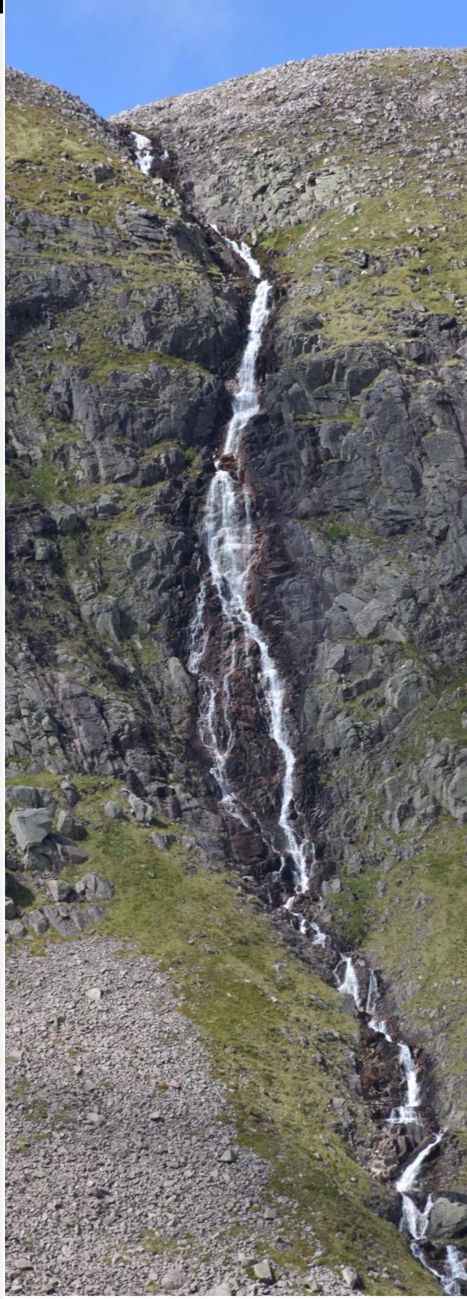

CALCULATING SHIFTS

A matrix expression for model shift directly from data shift – to first order

Answer: Yes, approximately.

- Analysis of Contingency Tables – see eg Christensen's book on Matrix Methods
- Adaptive Learning

We have taken another approach - ideas from Statistical Geometry – intuitive and capable of generalisation

Ingredients of the calculation

1. Design matrix for the log-linear model D
2. Metric Tensor at x and at y

$$G(x) = D^T \operatorname{diag}(x) D$$

3. Model shift equation

$$\delta y = \operatorname{diag}(y) D G(y)^{-1} D^T \delta x$$

		Shifted Data		Expected defaults on Shifted Data	
A1	A2	d	n	True Model	$y + \delta y$
1	1	60	1010	85.05	85.70
2	1	26	121	0.96	0.30
1	2	45	99	24.41	23.93
2	2	7	1002	27.59	28.05
1	3	13	95	8.52	8.35
2	3	4	1002	8.48	8.65

CALCULATING SHIFTS

A matrix expression for shifts of parameter shift and of other metrics – to first order

Change of model coefficients

$$\delta b = G(y)^{-1} D^T \delta x$$

	Int	A1=2	A2=2	A2=3
True Model on Shifted data	-3.164	1.224	-0.447	0.823
$b+\delta b$	-3.160	1.525	-0.919	1.034
Base Model = b	-3.731	2.250	-1.364	1.310

Change of model performance – for example PSI(observed, expected)

$$W = \sum (x - y) \log \frac{x}{y}$$

$$\delta W = (u + v D G(y)^{-1} D^T) \delta x$$

$$u = \text{row}\left(1 - \frac{y}{x} - \log \frac{y}{x}\right)$$

$$v = \text{row}\left(1 - \frac{x}{y} - \log \frac{x}{y}\right)$$

	PSI(observed,expected)
Base Model on base data = W	25.82
True Model on shifted data	151.44
$W+\delta W$	150.23
Base Model on shifted data	160.57

REBALANCING DATA

Compensating the known defects of data rebalancing

Data balancing is often used to improve model convergence and reduce errors. However such operations can distort the parameters of the underlying model.

In this example, default is rebalanced to be 50% overall – a massive disruption that this poorly specified model does not accommodate well.

But we can restore (some of) the true model's parameters by applying a data shift that reverses the rebalancing.

A1	A2	Z	Base	Balanced	δx
1	1	0	950	950	0
1	1	1	50	1200	-1150
2	1	0	95	95	0
2	1	1	5	120	-115
1	2	0	50	50	0
1	2	1	50	1200	-1150
2	2	0	995	995	0
2	2	1	5	120	-115
1	3	0	80	80	0
1	3	1	20	480	-460
2	3	0	998	998	0
2	3	1	2	48	-46

	Int	A1=1	A2=1	A2=2
Model on balanced data	-0.550	1.874	-0.876	0.873
Model shifted from balanced	-12.089	2.235	-1.256	1.391
Base = true model	-3.731	2.250	-1.364	1.310

MONITORING MODEL SHIFT

We can tell when PSI is in fact an effective metric

Suppose we're interested in the overall change of model caused by the data shift.
Choose an information metric to measure shifts: PSI(original model, shifted model)

The key matrix in this analysis is $M = G(y)^{-1}G(x)$ and its eigenvalues

x Data Observed Counts A2

	d/n	1	2	3
A1	1	20 / 30	10 / 15	30 / 50
	2	25 / 40	12 / 20	35 /60

Eigenvalues of M

1.034 1.023 1.015 1.010 1.000 1.000 0.990 0.984 0.980 0.961

With this spread of eigenvalues, PSI of (tangential) data shift is an effective metric

(0.961)* PSI(TDS) <= PSI(model shift) <= (1.034)*PSI(TDS)

Where TDS = tangential data shift = $\mathrm{diag}(x)DG(x)^{-1}D^T\,\delta x$

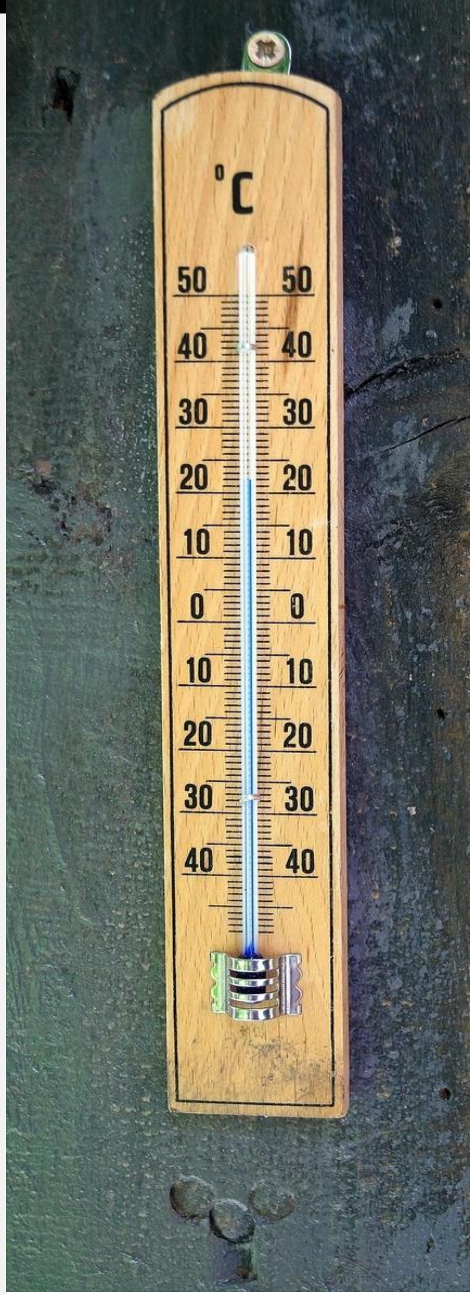

EARLY WARNING MONITORING

Early Warning Monitoring looks at model shifts caused only by population shifts

We observe only population shifts long before the model outcome, and we would like to know quickly how the population shift could damage the model.

We need Early Warning Monitoring and this can be set up directly in our model shift framework.

Start with the matrix H that converts data shifts into shifts of the outcome (y, b, W, etc) we're interested in. Build a projection matrix P that singles out population movements – based on a block decomposition between population shifts and parameter shifts

$$G(x) = \begin{pmatrix} A & C \\ C^T & B \end{pmatrix} \qquad P = \mathrm{diag}(x) D \begin{pmatrix} I & A^{-1}C \\ 0 & 0 \end{pmatrix} G(x)^{-1} D^T$$

Now work with HP instead of H.

Using the singular value decomposition of HP we find the movements of population that cause the largest movements to the outcome.

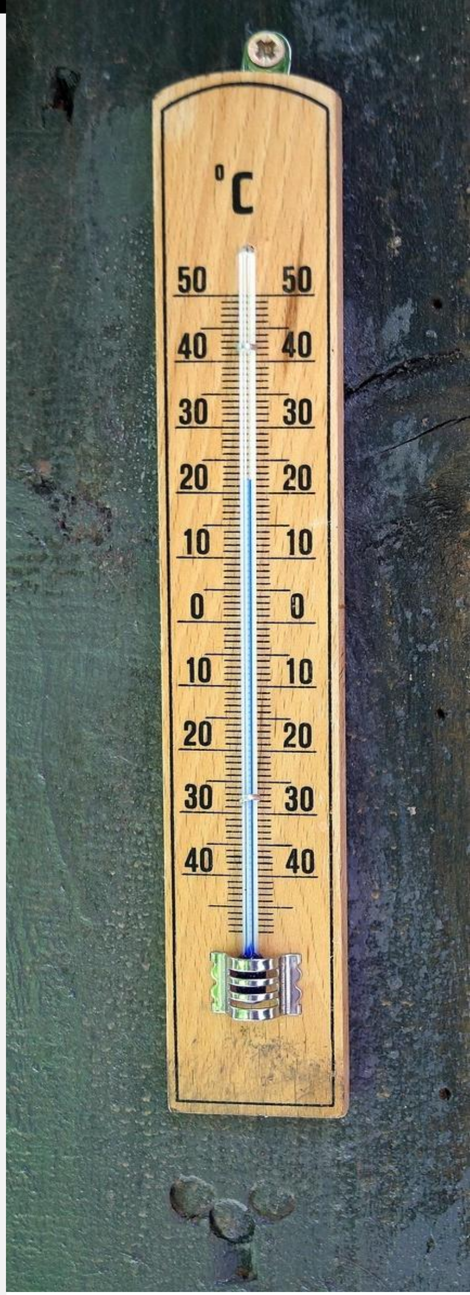

MONITORING MODEL SHIFT

Replace PSI with a weighted sum of data shifts – a model risk scorecard

For the imbalanced example before – no surprise that the eigenvalues of M are too well spread for tight constraints on PSI

Eigenvalues

1.499 1.045 1.012 1.004 1.000 1.000 0.978 0.940 0.911 0.743

Here PSI is better replaced by a different kind of metric – a weighted sum of data shift values

We wish to detect when model parameters will shift far from their original values.
For early detection we wish to use only the population shift.

We pick the weights using the singular values and vectors of the Singular Value Decomposition (SVD) of the δb matrix, with population projection restriction for an early warning version

$$G(y)^{-1} D^T P$$

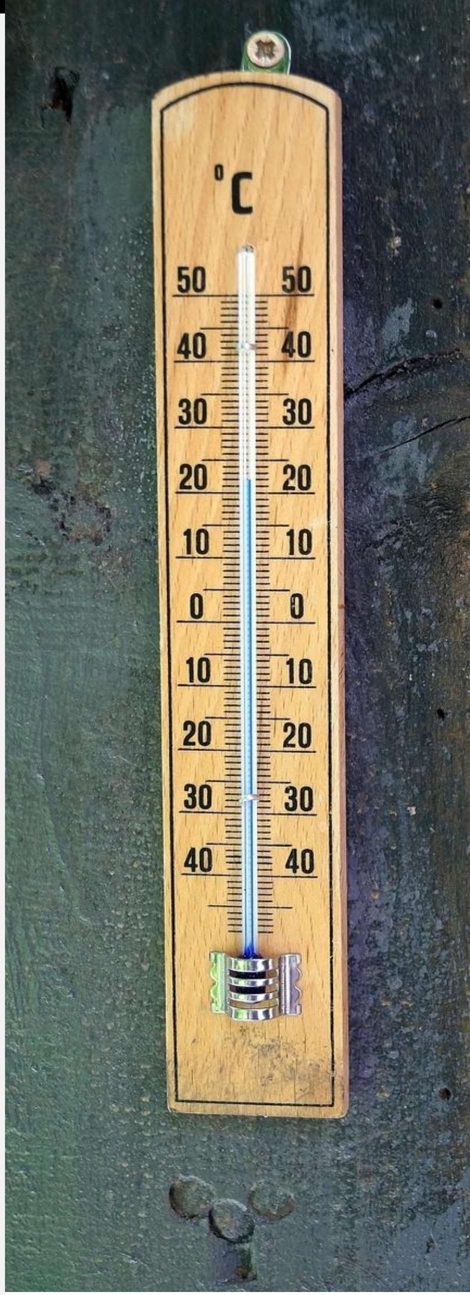

MONITORING MODEL SHIFT

A model risk scorecard in practice

The top singular value gives the following monitoring metric:

1. SVD gives weights for population shifts that give largest impact on model parameters

A1,A2	1,1	2,1	1,2	2,2	1,3	2,3
weights*1000	0.082	-9.974	0.286	4.237	-0.137	8.829

2. Take δx (population movement only) and score it up using those weights

δx	10	21	-1	2	-5	2

Score
-0.182

3. Check whether this score has exceeded an agreed threshold – via impact

	Int	A1=1	A2=1	A2=2
Impact on model when score = 1	5.708	-5.284	-0.736	0.759

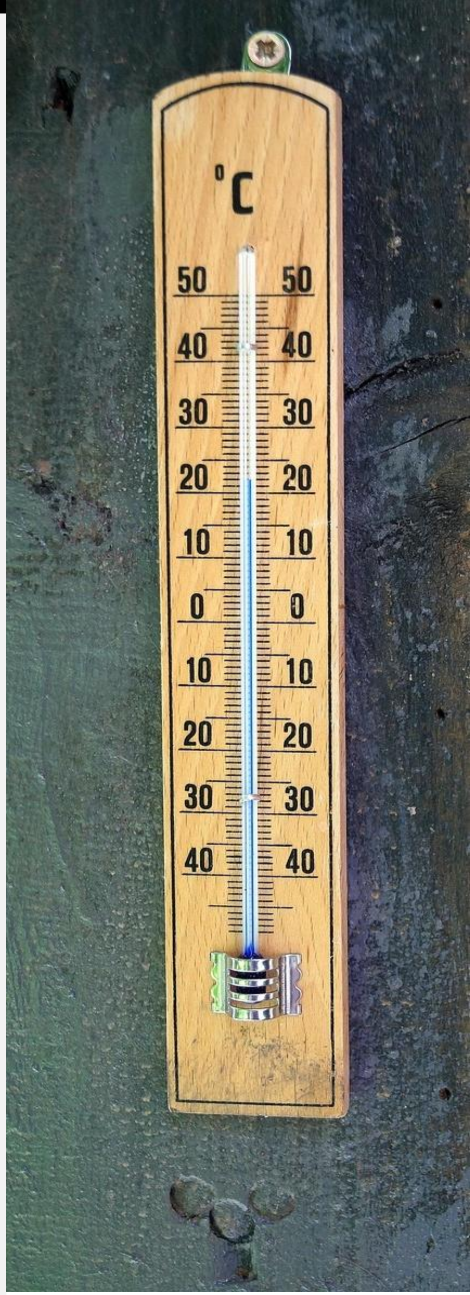

CONCLUSIONS

- Data Drift is not always damaging, and it is poorly monitored by blunt metrics like PSI.

- Understood and measured in the right way, Data Drift can be quantified and split into benign components and damaging components.

    - A tool to analyse and challenge models by applying challenge scenarios – Data Shift and Model Shift ;
        - All examples in this presentation are derived from R code under development.

    - Practical monitoring and early warning of damaging Data Drifts;
        - the specification of this monitoring is generated automatically at model development.
        - all model developments should be accompanied by their model risk monitoring specification.

- Let’s get good with Data Shifts – in validation and in model risk monitoring.