Efficiently exploring indexation in pension expenditure forecasts

Roman Goebel · Keva · 33rd International Congress of Actuaries, Tokyo 2026

© Roman Goebel, Keva

This presentation has been prepared for the 33rd International Congress of Actuaries.
The Institute of Actuaries of Japan (IAJ) and the International Actuarial Association (IAA) assume no responsibility for the opinions or content expressed in these presentation slides. Furthermore, the opinions and content contained herein do not necessarily reflect the views of the IAJ or the IAA.

About this work

Roman Goebel ·Head of Actuarial Affairs, Keva

1

ORIGIN

A rewrite of the author's actuarial qualification thesis, expanded from practical development work at Keva.

2

INTERNATIONAL REWRITE

Prepared for an international actuarial audience; Finnish system details remain only where they explain the method.

3

EXTENDED SCOPE

Adds stronger links to an ALM framework and recent results from production model tests.

Practical aim: explore many index paths without repeatedly running the production model.

AI disclosure: Generative AI supported the design of this presentation; all scientific content, calculations, interpretations and conclusions in the paper remain the author's responsibility.

Uniform benefit rules—differences in financing

FINNISH EARNINGS-RELATED PENSIONS

Mandatory statutory provision; largely uniform for insured persons

PRIVATE SECTOR

One jointly liable private-sector scheme

- Pension insurance companies
- Employer pension foundations
- Industry-specific pension funds

Provision distributed across several institutions

scheme — — — — administrator

PUBLIC SECTOR

KEVA

single administrator

Several financially independent schemes

All administered by Keva

Different structures; common rules and indices.

Overview: four parts

1	**The forecasting challenge**	Why many index paths overwhelm repeated model runs.
2	**The inverse solution**	How basis runs reveal reusable expenditure components.
3	**Evidence and insight**	Accuracy, speed-up and hidden model structure.
4	**Applications, boundaries and further development**	Where the method adds value—and where it stops.

PART 1

The forecasting challenge

1

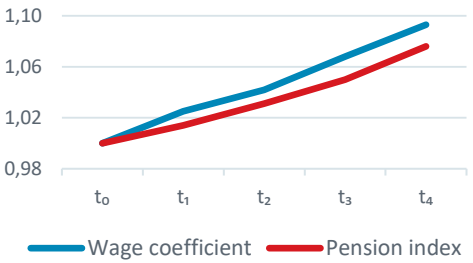
Why repeated production model runs constrain scenario exploration.

1 OF 4

A pension expenditure model is a program

INPUTS

Indexation



Population · earnings · legislation

PRODUCTION MODEL

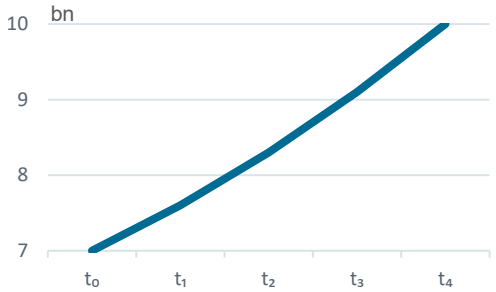
</>

PENSION EXPENDITURE MODEL

applies data, rules and assumptions year by year

OUTPUT

Annual expenditure E(t)



one value for every forecast year

New index curves → new model run → new expenditure path.

Brute-force reruns do not scale

A deterministic forecast

1 path

Production model run
+ expert operation

Stochastic and ALM analysis

thousands

of paths × changing outlooks
× dynamic decisions

The model becomes the bottleneck

One trusted model. Thousands of index paths.

1
tested forecast model

→

0
code changes

→

Thousands
of index paths

One-off basis runs—no production model rerun for each scenario.

Can indexation be explored externally while the tested forecast model remains authoritative?

PART 2

The inverse solution

2

Turn a finite set of basis runs into reusable expenditure components.

Move indexation outside the production model

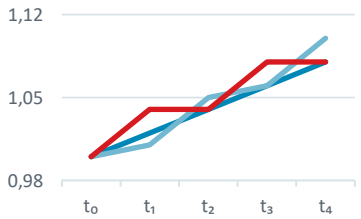


The production model remains the source of truth; only indexation is explored externally.

The inverse problem reveals hidden structure

1 KNOWN INPUTS

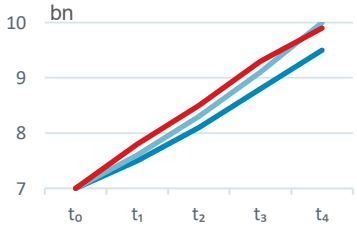
What I put in



Chosen index paths

2 OBSERVED OUTPUTS

What I get out



Observed expenditure paths

3 ASSUMED FORM

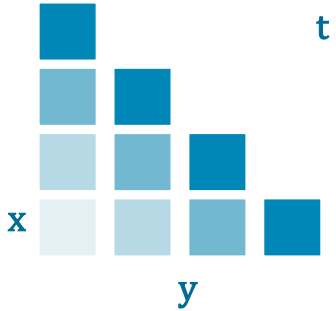
How they relate

$$E(t,I) = \Sigma\, E(x,y,t) \times r(x,y,t,I)$$

Unknown components × known factors

4 SOLVED PARTS

What I recover



Components E(x,y,t)

Σ unknown components × known factors = observed outputs
→ solve for the components

Both pension building blocks are sums of indexed earnings

OLD-AGE PENSION

Annual earnings over the working career + deferred rights accrued before 2005

$$E_y = \sum_{x=2005}^{y} a_x s_x \left(I_y^W / I_x^W\right) + E_{2004}\left(I_y^W / I_{2004}^W\right)$$

Each earnings year x contributes separately.

Every contribution carries the wage-coefficient ratio.

DISABILITY PENSION

In addition to the above, a projected pension component is added for the remaining period to retirement age *w*

$$\widehat{E}_y = \sum_{x=y-5}^{y-1} (w-y)\, a_y s_x \left(I_y^W / I_x^W\right)$$

Uses recent earnings and the remaining projected period.

The calculation repeats the same indexed-earnings structure.

Disability pension = target pension + projected component

The shared multiplicative index factor is what the inverse method exploits.

Pension indexation happens in two stages

BEFORE RETIREMENT

Recorded earnings — from the working career

Wage coefficient — updates accrued rights

Pension starts — the index regime changes

AFTER RETIREMENT

Pension in payment — monthly benefit

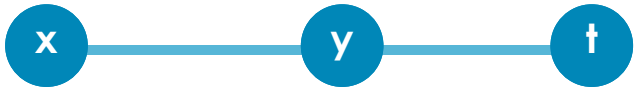
Pension index — annual adjustment

Expenditure E(t) — total paid in year t

Think of each payment as three dates

A COMPONENT E(x,y,t)

Accrual → pension start → payment



a right accrues — pension starts — expenditure paid

Each component is fixed once the model has been decoded.

THE FACTORISATION

$$E(t,I) = \Sigma\ E(x,y,t) \times r(x,y,t,I)$$

fixed component × indexation factor

Plain language: rebuild total expenditure by re-indexing and summing the components.

Recover the components in three steps

1

RUN A BASIS

Use single-year index increases
to create linearly independent runs.

2

SOLVE THE INVERSE

Combine model outputs with simple differences to recover E(x,y,t).

3

VALIDATE

Compare external reconstruction
with fresh production model runs.

Practical requirement: the model must allow index assumptions to be changed freely.

PART 3

Evidence and insight

3

Test the approximation, quantify the speed-up, and reveal hidden structure.

After setup, scenarios become cheap

≈9 hrs

for 5,151 basis paths

one-off production model extraction, 100 years, **nine** processor cores

≈6 min

for 66,000 paths

on **one** processor core

≈10,000

scenarios / minute

after the basis exists

The same compute budget could support ≈50 million scenarios.

Basis creation: 540 minutes on 9 cores. The payoff comes from repeated reuse.

Tests on production models show high accuracy

15-YEAR PRODUCTION MODEL TEST

Exact to the cent

2021–2035, including earnings development

100-YEAR PRODUCTION MODEL TEST

+0.03%
indexation only

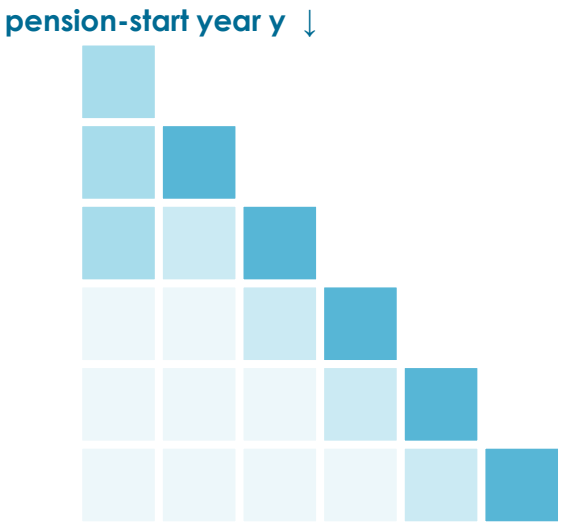
−0.02%
with earnings development

Differences remain small over 2025–2124

MODEL A ≠ **MODEL B**

Fundamentally different production models—not one model with two forecast horizons.

The components reveal how the model behaves



Conceptual component map—not to scale

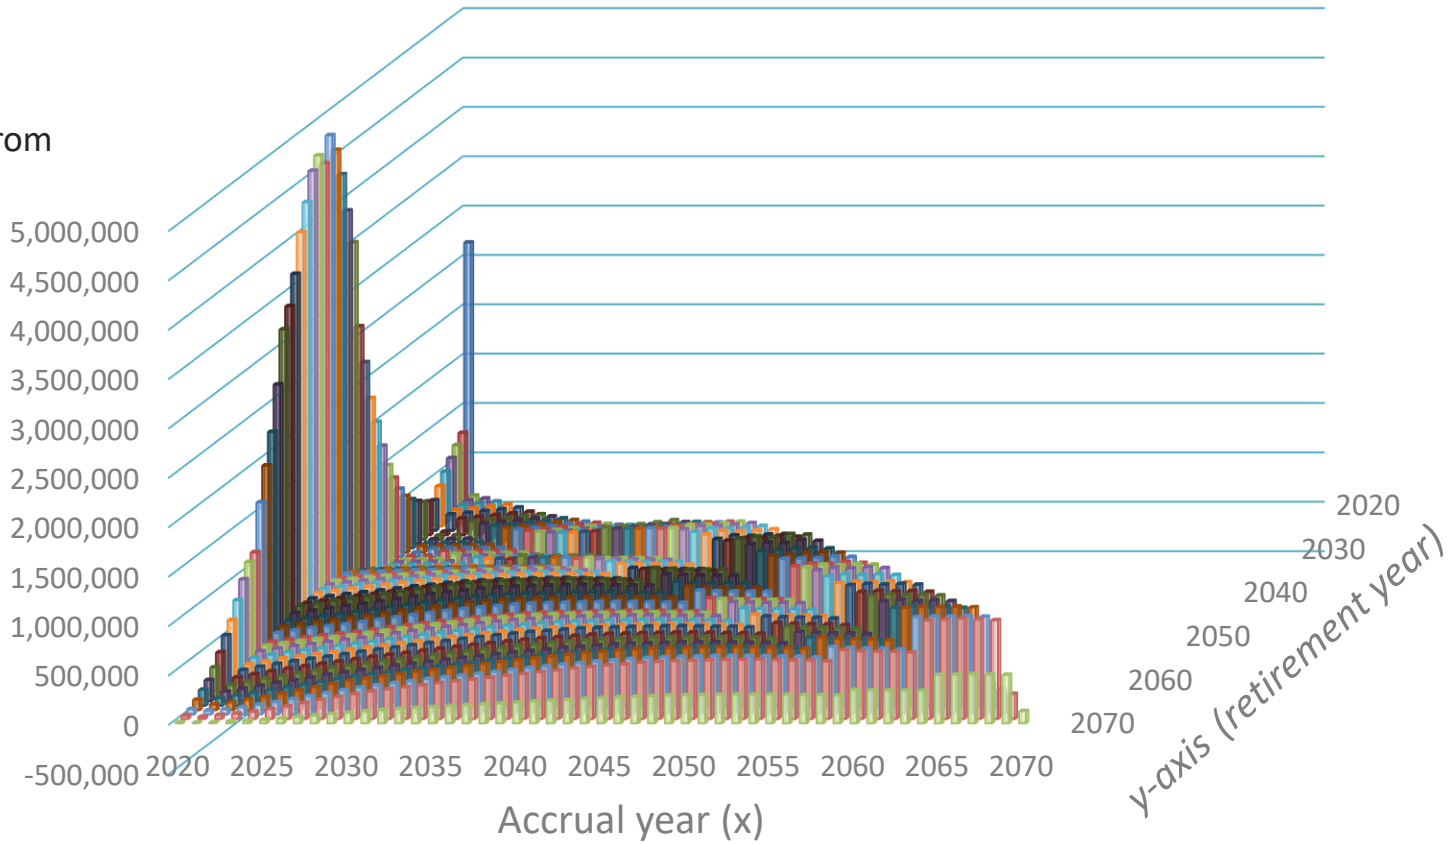
- Existing pension stock remains visible
- Future accrual flows into later payments
- Model changes can be compared component by component

UNINDEXED COMPONENTS IN 2070

Artificial example from the paper

E(x,y,2070)

Total: EUR 715.9 million



Public-sector tests support wider Finnish applicability

1

SYSTEM-WIDE INDICES

The wage coefficient and pension index
are common across the system.

2

SAME CORE STRUCTURE

Accrual → commencement → payment
follows the same statutory logic.

3

VALIDATE EACH MODEL

Private-sector models still require
independent validation.

PART 4

Applications, boundaries and further development

4

Use the representation where it adds value—and know where it stops.

The same representation supports four uses

OPERATIONAL

1 Fast forecast updates

Swap index assumptions without a full forecast rerun.

2 Model and policy review

Compare model versions and trace the intergenerational impacts of legislative changes.

STRATEGIC

3 Stochastic expenditure

Simulate many wage and price inflation paths.

4 Dynamic ALM

Connect expenditure, contributions and risk decisions.

Powerful—but only inside clear boundaries

REQUIRES

- Indexation assumptions that can be varied freely
- Basis runs within the model's valid range
- Independent validation for each forecasting model

DOES NOT

- Replace the pension expenditure forecasting model
- Generate wage, price or return scenarios
- Generalise automatically beyond the validated model

The next step: a dynamic ALM loop

Forward-looking assumptions change—and contribution and risk decisions react.

1

UPDATE THE OUTLOOK

New wage, price and return views reshape the paths still ahead.

2

RECALCULATE THE SYSTEM

Lightweight expenditure feeds the sustainable contribution level and required return.

3

DECIDE—THEN REPEAT

Contribution and allocation choices affect the next state—and the next outlook.

Efficient pension expenditure calculation leaves scenario generation as the main challenge.

Inverse problems deserve a wider actuarial role

The pension-index case is one concrete example of a wider framework.

1	**QUESTION**	Can the effect of the parameter be expressed in a factorised or otherwise separable form?
2	**TEST**	Can a finite, informative basis of production model runs identify the hidden components?
3	**VALIDATE**	Do the recovered components reproduce independent model runs inside the intended domain?

IF ALL THREE HOLD — **Keep the production model authoritative—reuse the recovered structure for scenarios, review and analysis.**

Keep the model. Make exploration lightweight.

1

Preserve trust

The tested production model remains authoritative. No code changes.

2

Extract structure

A finite basis reveals reusable expenditure components.

3

Scale insight

Fast scenarios enable stochastic analysis and ALM.

Thank you ·Questions