

The Workforce Claims Forecast.

Experience Calibration: the renewal math a carrier already runs on the employer's population, run by the buyer side — the group's own paid medical and Rx experience, credibility-blended against a manual rate, with Metra's forward metabolic signal folded in as a disclosed, auditable dollar offset, returned as a projected forward claims figure with a confidence band.

Metra Healthcare Intelligence · Actuarial Brief v2.0 · Experience Calibration addendum ·

Published with full assumption disclosure

This brief sets out the quantitative basis for Metra's figure. It is the companion to the v1.0 Actuarial Brief, which establishes how Metra measures workforce metabolic movement and monetizes it into a forward exposure figure. v2.0 takes that finished exposure figure and does with it what a carrier already does at renewal: it nets it against a credibility-blended projection of the group's own paid claims.

The central design principle carries forward from v1.0 and is sharpened here: the projection is anchored in **two quantities the group actually owns** — its real paid experience and its real measured metabolic movement — blended against a disclosed manual rate. Nothing in the pipeline is a proprietary black box, and the two credibility weightings in the system are kept rigorously distinct.

HOW TO READ THIS BRIEF — CONNECT, CREDIBILIZE, PROJECT

Connect: the group's own aggregate paid medical and Rx experience is matured to an incurred basis and large claimants are pooled out. **Credibilize:** the resulting observed per-employee-per-year (PEPY) rate is blended against the carrier manual rate using Bühlmann credibility, $Z = N/(N+400)$. **Project:** the blended rate is carried forward one year under trend and attrition, then reduced by the forward metabolic dollar signal from v1.0 — returning a projected claims figure with a confidence band.

ADDITIVE TO V1.0 — NOT A REPLACEMENT

The v1.0 Workforce Exposure Forecast and its published [Methodology v1.0](#) are unchanged and remain the version-pinned specification for the metabolic exposure figure. v2.0 consumes that figure as a finished input and is specified in full at [Methodology v2.0](#). This brief is the narrative companion to that specification; where the two differ on a number, the methodology document governs.

The Renewal Math, Run Buyer-Side

Carriers, PBMs, and stop-loss writers have run forward claims projections on employer populations for decades: take the group's experience, credibility-weight it against a manual rate, trend it forward, and price the renewal off the result. The employer has never run that arithmetic for itself. The Workforce Claims Forecast is the buyer-side counterpart — the same renewal math, computed by the employer, with one disclosed addition the carrier does not have: a forward metabolic signal drawn from the workforce's own measured movement.

"v2.0 does not invent a new actuarial method. It runs the renewal blend the carrier already runs, on the group's own experience, and folds Metra's measured metabolic offset into it as a transparent line item the actuary across the table can check and remove."

Anchored in the group's own experience

The observed rate is the group's actual paid claims, matured and pooled — not a benchmark and not an assumption. Where the group supplies its carrier manual rate, the entire blend rests on the group's own numbers.

Credibility-weighted, not asserted

The observed rate is blended against the manual rate by a stated Bühlmann weight. A small group leans on the manual; a large group leans on its own experience. The weight is reported with every figure.

The metabolic offset is a disclosed line item

Metra's forward exposure figure enters as a single subtracted quantity with its own confidence interval. A reviewer who rejects it can set it to zero and read the gross projection underneath.

Reproducible from disclosed inputs

The engine is pure and deterministic and carries a reproducibility digest. Every signed figure can be regenerated from the inputs recorded against it.

Three Layers: Connect, Credibilize, Project

The calibration separates into three layers. Keeping them distinct is what lets the projection rest on owned quantities while the manual rate and the metabolic offset each do the work they are suited for — anchoring the small-group case and crediting measured movement, respectively.

LAYER 01 — CONNECT

The group's aggregate paid medical and Rx for a defined period, with the enrolled member-months that generated them, are connected as a single record carrying a source label and a paid-through date. Paid is matured to an incurred basis by a supplied completion factor (preferred) or an explicit IBNR add-on. Dollars above a supplied large-claim pooling point are excluded from the credibility base, so a single catastrophic claimant cannot distort the group's rate. No member-level record, diagnosis, or identifier is involved — the layer is aggregate-only by construction.

LAYER 02 — CREDIBILIZE

The pooled, completed base per average enrolled life is the observed PEPY. It is blended against the manual / expected PEPY by the Bühlmann weight $Z = N/(N+400)$, where N is average enrolled lives. The blend is the standard renewal one: at $N = 400$ the weight is exactly one-half; below it the forward rate leans on the manual, above it on the group's own experience.

LAYER 03 — PROJECT

The blended rate is carried forward onto the surviving lives (after attrition), trended one year, and then reduced by the v1.0 metabolic offset. The result is the projected forward claims figure. The confidence band on the figure is inherited from the metabolic offset's confidence interval, inverted so that a larger avoided-exposure offset maps to a lower projected-claims figure.

Maturing and Pooling the Connected Experience

Paid claims for a recent period are immature — claims incurred near period end are not yet paid. The procedure matures paid to an incurred basis in one of two mutually exclusive ways, in priority order, both supplied by the actuary rather than assumed by the platform.

COMPLETION FACTOR (PREFERRED)

$$C = P / c$$

EXPLICIT IBNR ADD-ON (ONLY WHEN NO COMPLETION FACTOR IS SUPPLIED)

$$C = P + R$$

Where a pooling point is supplied, the dollars above it are removed from the credibility base, and the excluded amount is reported so a reviewer can reconstruct the unpooled figure.

POOLED CREDIBILITY BASE

$$B = \max(0, C - X)$$

The completion factor and the pooling point are the actuary's levers. Neither is estimated from the data; both are disclosed inputs that move the observed rate, and the output reports what was supplied versus defaulted. Where the group has not supplied a manual rate, the expected term falls back to a generic national benchmark, which is explicitly labelled as a placeholder to be replaced — a benchmark-anchored blend is materially weaker than a manual-anchored one, and the document says so.

The Credibility Blend

The observed PEPY is the pooled, completed base per average enrolled life. It is blended against the manual / expected PEPY by the Bühlmann credibility weight, using the same credibility constant family as the v1.0 forecast so a reviewer meets one credibility parameter across the whole instrument.

OBSERVED PEPY

$$O = B / N, \quad N = \text{member-months} / 12$$

CREDIBILITY WEIGHT

$$Z = N / (N + 400)$$

BLENDED PEPY

$$\Pi = Z \cdot O + (1 - Z) \cdot E$$

GROUP SIZE (N)	CREDIBILITY Z	FORWARD RATE LEANS ON
100 lives	0.20	Mostly the manual rate
400 lives	0.50	Evenly split
1,000 lives	0.71	Mostly the group's own experience
2,000 lives	0.83	Predominantly own experience

Forward Projection and the Metabolic Offset

The blended rate is carried forward onto the surviving lives, trended one year, then reduced by the forward metabolic dollar signal from the v1.0 Workforce Exposure Forecast.

FORWARD LIVES

$$L = N \cdot (1 - \alpha)$$

GROSS PROJECTED CLAIMS

$$G = L \cdot \Pi \cdot (1 + t)$$

PROJECTED FORWARD CLAIMS

$$\hat{G} = G - M$$

Here M is the v1.0 12-month point estimate — the forward healthcare exposure the cohort's measured metabolic trajectory is expected to avoid over the year — entering as finished dollars. The trend t defaults to a conservative single-digit anchor and is intended to be overridden with the carrier's renewal trend; the attrition α defaults to the v1.0 value. The implied trend impact of the offset, M / G , is reported so the offset can be read as a fraction of gross.

Two Credibility Weights, Never Stacked

The single most important integrity property of the calibration — and the first thing a reviewing actuary should verify — is that the system contains **two distinct credibility applications, applied to two different quantities, that are never compounded.**

WEIGHT	CREDIBILIZES	AGAINST	WHERE IT LIVES
Z_{claims}	The group's observed paid PEPY	The manual / expected PEPY	This brief, Layer 02
$Z_{\text{metabolic}}$	The cohort biometric exposure signal	A conservative prior	v1.0, inside the offset M

The metabolic offset M arrives already credibility-weighted by v1.0's $Z_{\text{metabolic}}$. v2.0 consumes it as a fixed dollar figure and applies **no further credibility weight to it**. Z_{claims} touches only the observed-versus-expected PEPY blend and never touches M. The two weights operate on disjoint quantities; there is no path by which a single dollar of exposure is credibility-discounted twice.

REFERENCE CASE

322-Life Reference Case: Calibration Mechanics

A 322-life group with a full plan year of connected experience, a 0.92 completion factor, a \$150,000 pooling point with \$380,000 above it, a carrier manual rate of \$16,500 PEPY, an 8% forward trend, 15% attrition, and a v1.0 metabolic offset of \$240,000 (CI \$160,000–\$330,000).

STEP	CALCULATION	RESULT
Average enrolled lives	$3,864 \text{ member-months} \div 12$	N = 322
Mature paid → incurred	$\$5,300,000 \div 0.92$	\$5,760,870
Pool large claimants	$\$5,760,870 - \$380,000$	B = \$5,380,870
Observed PEPY	$\$5,380,870 \div 322$	\$16,711
Credibility weight	$322 \div (322 + 400)$	Z = 0.446
Blended PEPY	$0.446 \times 16,711 + 0.554 \times 16,500$	$\pi = \\$16,594$
Gross projected	$273.7 \text{ lives} \times 16,594 \times 1.08$	G = \$4,905,121
Less metabolic offset	$\$4,905,121 - \$240,000$	$\hat{G} = \\$4,665,121$
Confidence band	G – offset CI, inverted & sorted	\$4,575,121 – \$4,745,121

Result. Projected forward claims of **\$4,665,121**, a 95% band of **\$4,575,121 – \$4,745,121**, at a claims credibility of 44.6% on the group's own experience. The metabolic offset represents 4.9% of gross projected claims — a deliberately modest, defensible line item, not a headline savings claim.

How the Confidence Band Is Constructed

The band on the projected figure is inherited entirely from the v1.0 metabolic offset's confidence interval. The gross projection G is treated as a point quantity for this purpose; the band reflects the uncertainty in the modelled metabolic offset, which is the quantity Metra estimates.

LOWER / UPPER BOUND (SORTED SO LOW \leq HIGH)

$$\hat{G}_{lo} = G - M_{hi}, \quad \hat{G}_{hi} = G - M_{lo}$$

The mapping is deliberately inverted: a larger avoided-exposure offset produces a lower projected-claims figure, so the high end of the offset CI maps to the low end of the projected-claims band. The band does not attempt to layer a second source of statistical uncertainty onto the connected paid experience, which is the group's own and is reported as supplied. A reviewer who wishes to stress the trend or completion assumptions can do so directly, since both are disclosed inputs.

Aggregate-Only by Construction

No protected health information

The connected-experience record stores period-level dollar totals and member-months — never a member-level claim line, diagnosis, procedure, or identifier. The metabolic offset it consumes is itself a cohort-level figure subject to the v1.0 small-cell governance floor. The calibration never touches, joins, or re-identifies PHI, and no member's claim or biometric history can be reconstructed from the output.

Access control

Both the calibration inputs and the signed output are gated to the employer's forecast-authorized administrators. A non-authorized administrator cannot reach the connected experience or generate a forecast.

Audit and reproducibility

Every generated calibration is recorded with the inputs' provenance, the resulting credibility weight, the methodology version, and a reproducibility digest of the inputs — so any signed figure can be regenerated and independently checked.

Conformance with Actuarial Standards of Practice

ASOP	SUBJECT	HOW THE CALIBRATION CONFORMS
No. 23	Data Quality	Connected experience carries a source label and paid-through date; completion and IBNR are disclosed inputs; benchmark fallback is explicitly flagged as a placeholder.
No. 25	Credibility Procedures	Bühlmann weight $Z = N/(N+400)$ with a stated, common constant; observed, expected, and weight reported; the two credibility applications kept disjoint.
No. 41	Actuarial Communications	The signed output names the methodology version, signer tier, assumptions, and a reproducibility digest; an unsigned tier is labelled methodology-only.
No. 56	Modeling	The engine is pure, deterministic, unit-tested, and reproducible from its digest; assumptions are externalized and disclosed, not embedded.

What the Calibration Does Not Claim

The blend is only as good as the manual rate

When no group manual rate is supplied, the expected term falls back to a national benchmark and the output is flagged. A benchmark-anchored blend should not be carried to a renewal table without the carrier's own manual rate.

Completion and trend are supplied, not derived

The procedure applies the completion factor and trend the actuary supplies; it does not estimate them. A wrong completion factor moves the observed rate proportionally and a wrong trend moves the entire gross projection.

The band reflects metabolic uncertainty, not experience volatility

The confidence band is inherited from the v1.0 offset CI. It does not capture random fluctuation in the group's own paid experience; that must be stressed against the disclosed inputs directly.

Cohort floor and single-period scope

The calibration inherits v1.0's minimum-cohort requirement (≥ 15 enrolled lives) and refuses to generate below it. It calibrates a single connected period to a single forward year; it is not a multi-year reserve model and does not chain successive renewals.

Not a carrier rate filing

The Workforce Claims Forecast is a buyer-side instrument carried into the renewal or underwriting conversation alongside — never instead of — the carrier-side instruments that already exist. It does not certify a rate.

CITATION REGISTER

1. Bühlmann H. Experience rating and credibility. *ASTIN Bulletin*. 1967;4(3):199–207.
2. Klugman SA, Panjer HH, Willmot GE. **Loss Models: From Data to Decisions**. 5th ed. Wiley; 2019.
3. Actuarial Standards Board. **ASOP No. 23**, Data Quality. Actuarial Standards Board.
4. Actuarial Standards Board. **ASOP No. 25**, Credibility Procedures. Actuarial Standards Board.
5. Actuarial Standards Board. **ASOP No. 41**, Actuarial Communications. Actuarial Standards Board.
6. Actuarial Standards Board. **ASOP No. 56**, Modeling. Actuarial Standards Board.
7. KFF (Kaiser Family Foundation). Employer Health Benefits Survey 2024. Section 1: Cost of Health Insurance.
8. PricewaterhouseCoopers Health Research Institute. Medical Cost Trend: Behind the Numbers 2025. PwC HRI; 2024.
9. Society of Actuaries. Group health renewal and manual-rating practice notes.
10. Metra Healthcare Intelligence. **Methodology v2.0 — Experience Calibration**. usemetra.com/methodology/v2.0/. The metabolic-exposure derivation, per-marker coefficients, and bootstrap confidence-interval construction consumed as the offset M are specified in full in Methodology v1.0.

Metra Healthcare Intelligence · usemetra.com · Prepared for actuarial review · v2.0