PREPARED FOR ACTUARIAL REVIEW · V2.1

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 two labeled uncertainty bands: an expected-claims confidence band and a realized-paid predictive interval.

Metra Healthcare Intelligence · Actuarial Brief v2.1 · 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. The calibration 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.

WHAT CHANGED IN V2.1

v2.1 is additive to v2.0 and changes two things only. **(1) The credibility constant is now derived, not borrowed.** The value K = 400 is set as an independent limited-fluctuation full-credibility anchor (full credibility Z^{*} = 0.90 at N^{*} = 3,600 life-years ⇒ $K = N^{*}(1-Z^{*})/Z^{*}$ = 400); the worked numbers are unchanged. **(2) The single inherited band is replaced by two labeled bands** from one variance model — an expected-claims confidence band on the mean and a realized-paid predictive interval. v2.0 remains published and version-pinned at Methodology v2.0; attestations do not transfer across versions.

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 two labeled uncertainty bands.

ADDITIVE TO V1.0 AND V2.0 — 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. The calibration consumes that figure as a finished input and is specified in full at Methodology v2.1, which is itself additive to v2.0. This brief is the narrative companion to that specification; where the two differ on a number, the methodology document governs.

CORE CLAIM

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.

> *"The calibration 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 whose constant is derived from an explicit full-credibility standard. 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.

THE MODEL

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 and the constant 400 is a derived full-credibility anchor. 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, reported with two labeled uncertainty bands built from a single variance model: an expected-claims confidence band on the mean, and a wider realized-paid predictive interval.

LAYER 01 — CONNECT

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.

LAYER 02 — CREDIBILIZE

The Credibility Blend and Its Constant

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.

OBSERVED PEPY

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

CREDIBILITY WEIGHT

$Z = N / (N + k), \quad k = 400 \text{ life-years}$

BLENDED PEPY

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

Where the constant comes from (v2.1)

v2.0 carried the constant = 400 over from the v1.0 forecast as a convenience, so a reviewer met one credibility parameter across the instrument. v2.1 derives it independently instead, as a limited-fluctuation **full-credibility anchor**: fix the standard for full credibility at = 0.90, reached at = 3,600 life-years of exposure — a conventional group-health full-credibility threshold — and read off the Bühlmann constant that places the partial-credibility curve through that anchor.

FULL-CREDIBILITY ANCHOR → K

$k = N^* \cdot (1 - Z^*) / Z^* = 3{,}600 \cdot (1 - 0.90) / 0.90 = 400 \text{ life-years}$

So 400 is not a borrowed constant; it is the credibility constant implied by a stated full-credibility standard. The same value happens to govern v1.0's metabolic weight, but the two are now justified separately, on their own exposure bases — a reviewer may reset either anchor without touching the other.

Implied EPV / VHM — a reasonableness check

Bühlmann's constant is = EPV / VHM, the ratio of the expected process variance per life to the variance of the hypothetical means. Adopting the between-employer dispersion τ (the next section) as the VHM scale, = (τ·)², the constant implies a per-life process SD of = √( ) = √ ·τ· — an implied per-life pooled-claims CV of √ ·τ = √400·0.15 = **3.0**×. A per-life CV near 3× is entirely ordinary for one year of one person's pooled medical + Rx spend, where most lives spend little and a few spend a great deal. The check confirms the constant is not in conflict with the dispersion the band assumes; it is not an empirical estimate of EPV and VHM from member-level data, which the aggregate-only procedure never touches.

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
3,600 lives	**0.90**	Full-credibility anchor (Z∗)

LAYER 03 — PROJECT

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.

INTEGRITY PROPERTY

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
$Z_{metabolic}$	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 $Z_{metabolic}$. The calibration 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, the default between-employer dispersion $\tau = 0.15$, and a v1.0 metabolic offset of $240,000 (CI $160,000–$330,000).

STEP	CALCULATION	RESULT
Average enrolled lives	3,864 member-months ÷ 12	**N = 322**
Mature paid → incurred	$5,300,000 ÷ 0.92	**$5,760,870**
Pool large claimants	$5,760,870 − $380,000	**B = $5,380,870**
Observed PEPY	$5,380,870 ÷ 322	**$16,711**
Credibility weight	322 ÷ (322 + 400)	**Z = 0.446**
Blended PEPY	0.446 × 16,711 + 0.554 × 16,500	**π = $16,594**
Gross projected	273.7 lives × 16,594 × 1.08	**G = $4,905,121**
Less metabolic offset	$4,905,121 − $240,000	**$\hat{G}$ = $4,665,121**
Expected-claims band (95%)	$\hat{G}$ ± 1.95996 × $546,271	**$3,594,449 – $5,735,792**
Realized-paid predictive (95%)	$\hat{G}$ ± 1.95996 × $1,039,538	**$2,627,664 – $6,702,577**

Result. Projected forward claims of **$4,665,121**, with an expected-claims 95% confidence band of **$3,594,449 – $5,735,792** and a realized-paid 95% predictive interval of **$2,627,664 – $6,702,577**, 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.

UNCERTAINTY

Two Labeled Bands, One Variance Model

v2.0 inherited its band entirely from the v1.0 offset CI and treated the gross projection G as a point. v2.1 replaces that with two labeled bands built from a single Bühlmann–Straub variance model, so the figure carries both the uncertainty in the and the volatility of claims. The model has one disclosed dispersion input, τ — the between-employer CV of the true mean PEPY (default 0.15, overridable) — from which the VHM is = (τ·E)² and the per-life process SD is = √(k·a) = √k·τ·E.

(1) Expected-claims confidence band — the headline band

This is the estimation uncertainty of the forward claims, combining two independent sources in quadrature: the credibility estimation error of the blended PEPY, and the estimation error of the metabolic offset taken from its own 95% CI.

CREDIBILITY ESTIMATION SE OF THE BLENDED RATE, SCALED FORWARD

$$SE(\hat{\pi}) = \tau \cdot E \cdot \sqrt{(1 - Z)}; \quad SE(G) = L \cdot (1 + t) \cdot SE(\hat{\pi})$$

OFFSET SD FROM ITS 95% CI, THEN THE BAND

$$SD(M) = (M_{hi} - M_{lo}) / (2 \cdot 1.95996); \quad SD_{exp} = \sqrt{(SE(G)^2 + SD(M)^2)}; \quad \hat{G} \pm 1.95996 \cdot SD_{exp}$$

(2) Realized-paid predictive interval — labeled, pre-risk-charge

A renewal is not only an estimate of the mean — next year's pooled claims fluctuate around that mean. The predictive interval adds the forward-year process variance of the retained block to the estimation band. It is labeled separately and reported pre-risk-charge: it is the statistical spread of paid claims, not a priced premium, and it excludes the stop-loss layer above the pooling point.

FORWARD-YEAR PROCESS SD OF THE POOLED BLOCK, THEN THE PREDICTIVE INTERVAL

$$SD_{proc} = (1 + t) \cdot s \cdot \sqrt{L}; \quad SD_{pred} = \sqrt{(SD_{exp}{}^2 + SD_{proc}{}^2)}; \quad \hat{G} \pm 1.95996 \cdot SD_{pred}$$

Both bands are centred on the projected figure Ĝ and are symmetric. The expected-claims band answers “how well do we know the expected cost”; the predictive interval answers “how wide could the realized paid claims land,” and is always the wider of the two. A reviewer who

wishes to stress trend, completion, or dispersion can do so directly through the disclosed inputs t, c, and τ.

DATA GOVERNANCE & PRIVACY

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. The dispersion τ is a disclosed assumption, not an empirical member-level estimate.

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.

STANDARDS CONFORMANCE

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 constant derived from an explicit full-credibility standard; 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, both uncertainty bands, 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.

MODEL LIMITATIONS

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.

Experience volatility is now in the band — within stated scope

v2.0's single band reflected only metabolic uncertainty. v2.1 adds an expected-claims confidence band and a realized-paid predictive interval, so experience volatility is now quantified through τ. What the predictive interval still excludes is stated, not hidden: it is pre-risk-charge, it does not re-add the stop-loss layer above the pooling point, and it does not price a carrier margin.

The metabolic offset is inherited from v1.0 and may double-count across markers

v1.0 monetizes weight, blood pressure, blood glucose, and waist on additive per-marker coefficients. When one physiological improvement moves several markers at once, the additive offset can overstate avoided exposure. Because the offset projected claims, the bias is conservative against the buyer — it makes the saving look larger, never the cost. v2.1 discloses this rather than silently correcting it; the planned fix is a v1.1 exposure haircut for cross-marker correlation, deferred to its own version bump and re-attestation.

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. Bühlmann H, Straub E. Glaubwürdigkeit für Schadensätze. **Bulletin of the Swiss Association of Actuaries.** 1970.
3. Longley-Cook LH. An introduction to credibility theory. **Proceedings of the Casualty Actuarial Society.** 1962 (limited-fluctuation full-credibility standards).
4. Herzog TN. **Introduction to Credibility Theory.** ACTEX Publications.
5. Klugman SA, Panjer HH, Willmot GE. **Loss Models: From Data to Decisions.** 5th ed. Wiley; 2019.
6. Actuarial Standards Board. **ASOP No. 23**, Data Quality. Actuarial Standards Board.
7. Actuarial Standards Board. **ASOP No. 25**, Credibility Procedures. Actuarial Standards Board.
8. Actuarial Standards Board. **ASOP No. 41**, Actuarial Communications. Actuarial Standards Board.
9. Actuarial Standards Board. **ASOP No. 56**, Modeling. Actuarial Standards Board.
10. KFF (Kaiser Family Foundation). Employer Health Benefits Survey 2024. Section 1: Cost of Health Insurance.
11. PricewaterhouseCoopers Health Research Institute. Medical Cost Trend: Behind the Numbers 2025. PwC HRI; 2024.
12. Society of Actuaries. Group health renewal and manual-rating practice notes.
13. Metra Healthcare Intelligence. **Methodology v2.1 — Experience Calibration.** usemetra.com/methodology/v2.1/. Additive to 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.