METHODOLOGY · V1.0 · FOR REVIEW & ATTESTATION

METRA Workforce Exposure Forecast — Methodology v1.0

The mathematical and actuarial methodology used to produce the per-client Workforce Exposure Forecast generated from the Metra platform. Published in the open, in v1.0, for review and attestation by credentialed actuaries.

Version 1.0 · **Published** May 2026 · **Status** For review & attestation ·

Permalink usemetra.com/methodology/v1.0/

Document version	Methodology hash	Bühlmann constant k	Lag haircut λ	Attrition α
1.0	**v1.0-2026.05**	**400 person-quarters**	**0.50**	**0.15 / yr**

Bootstrap iterations B
10,000

Foreword. This document specifies the cohort construction, credibility weighting, per-marker exposure estimation, confidence interval construction, and forecast horizon math used by the Metra platform to generate each client's Workforce Exposure Forecast. It is published openly in version 1.0 for review and attestation by credentialed actuaries (ASA, FSA, MAAA, FCAS, ACAS). Attestations are recorded at usemetra.com/methodology/attestations/. The intended primary readership is the carrier-side or independent actuary reviewing this document in the context of a renewal or stop-loss conversation; the document is written for that audience's bar, not for the marketing reader. Critique, correction, and counter-derivation are welcomed and will inform v1.1.

Posture. Every reference cited in this document is publicly verifiable from the reviewing actuary's own desk. Metra does not rely on proprietary internal claim tables, undisclosed coefficient derivations, or unpublished population priors. The complete methodology is in this document. There is no second document held back, no “commercial methodology” behind a contract. The companion at usemetra.com/methodology/companion/ is a reading guide, not a supplement.

CONTENTS

§1. Scope and Intent

§2. Notation

§3. Cohort Construction and Exposure Units

§4. Credibility Weighting (Bühlmann)

§5. Per-Marker Exposure Estimation and Confidence Interval Construction

§6. Forecast Horizon Construction

§7. METRA Score Composition

§8. Biometric-to-Claim Pathway and Lag Literature

§9. Per-Marker Cost Coefficients

§10. Aggregation to Workforce Total

§11. Limitations and Known Constraints

§12. Attestation Protocol

§13. Version History and Change Control

§14. Bibliography

§1. Scope and Intent

What this document specifies. The end-to-end mathematical procedure by which Metra produces, for a single enrolled employer cohort, a forward-looking estimate of healthcare-cost exposure attributable to the cohort's measured biometric trajectory. The estimate is returned at four horizons — 30, 60, and 90 days, and 12 months — with explicit confidence intervals and an explicit credibility weight on the cohort-specific signal relative to a conservative prior.

What this document does not specify. It does not specify a replacement for carrier-side rate filings, pooled experience tables, or stop-loss credibility methodologies. The Workforce Exposure Forecast is a complementary, buyer-side instrument intended to be carried into the renewal or underwriting conversation alongside — never instead of — the carrier-side instruments that already exist. Where this document discusses credibility, it does so in the Bühlmann sense applied to the cohort's biometric signal, not in the sense of carrier-side experience credibility on incurred claims.

Who is served by the output. The Forecast is generated for and read by four roles: the CFO planning against next year's exposure line; the CHRO governing the workforce population that drives the line; the broker carrying the document into the carrier conversation; and the actuary (employer-retained, broker-retained, or carrier-retained) folding the signal into the credibility math they already do.

Why this is published in the open. The Forecast is intended to be defensible at the renewal table. A defense requires the methodology to be auditable by the actuary across the table. The doctrine of openness is therefore not an aesthetic choice — it is the operational requirement that makes the instrument usable. Every coefficient, every formula, every assumption disclosed in this document is intended to be checked.

§2. Notation

The following symbols are used throughout this document. Where a symbol is reused in a section-specific sense, the local definition takes precedence and is noted inline.

N	Cohort size. The count of continuously enrolled employees in the observation window.
n	Exposure in person-quarters (PQ). One PQ = one employee with at least one valid reading for the marker during a calendar quarter. See §3.
k	Bühlmann credibility constant, expressed in person-quarters. v1.0 value: **k = 400 PQ**. See §4.
Z	Bühlmann credibility weight, $Z = n / (n + k)$, applied to the cohort-specific observed delta. $1 - Z$ is applied to the prior.
m	Index over the four monetized markers: weight, systolic BP, fasting glucose, waist circumference.
Δ_m	Cohort-mean per-employee improvement in marker m over the observation window, expressed as a positive number for improvement.
c_m	Per-unit annual cost coefficient for marker m, in USD per person-year per unit of sustained improvement. See §9.
E_m	Per-marker annualized exposure in USD attributable to the cohort's measured improvement in marker m.
E_{total}	Sum of per-marker exposures across the four markers.
h	Forecast horizon, expressed in days from the observation date (30, 60, 90, or 365).
λ	Lag haircut applied at the 12-month forecast horizon. v1.0 value: **λ = 0.50**. See §6.
α	Annualized workforce attrition rate, applied as a uniform decay. v1.0 value: **α = 0.15**. See §6.
B	Number of bootstrap iterations used for confidence-interval construction. v1.0 value: **B = 10,000**. See §5.

§3. Cohort Construction and Exposure Units

The cohort is the set of employees enrolled in the Metra platform under the employer's instance during the observation window. The window is defined as the calendar period from the employer's program start date to the Forecast generation date, with a maximum lookback of 24 months. Within this window:

- **Continuous enrollment** is required: employees who deactivated during the window are excluded from the active cohort but their pre-deactivation observations are retained for attrition estimation.
- **Marker-specific validity**: for each of the four markers (weight, systolic BP, fasting glucose, waist circumference), an employee contributes to that marker's cohort signal only if they have at least one valid reading in each calendar quarter of the observation window. Employees with sparse coverage on a marker contribute zero person-quarters for that marker.
- **Cohort floor**: company-wide aggregations require N ≥ 15; scoped aggregations (department, division, location) require N ≥ 5. Below floor, the Forecast does not produce a per-marker estimate for the scope.

Exposure unit: the person-quarter. A person-quarter (PQ) is the unit of validity-weighted observation used throughout the credibility and CI math. One employee continuously enrolled for the full 12-month observation window with at least one valid reading per quarter contributes 4 PQs of exposure on that marker. Total cohort exposure in PQs is:

> EXPOSURE IN PERSON-QUARTERS
>
> $$n_m = \Sigma_{i \in cohort} q_{i,m}$$
>
> *where $q_{i,m}$ is the count of calendar quarters in the observation window during which employee i has at least one valid reading for marker m.*

The choice of person-quarters as the exposure unit is deliberate. It matches the empirical cadence at which workforce-cohort biometric drift is reliably measurable — sub-quarter aggregation amplifies test-retest noise; supra-quarter aggregation hides meaningful drift. The Bühlmann credibility constant is expressed in the same unit (§4), and the worked examples that follow use this unit consistently.

Prior trajectory. For credibility weighting in §4, the no-intervention prior used in v1.0 is — no biometric movement. This is intentionally conservative. The published NHANES trends for working-age U.S. adults show a small annual deterioration in weight, fasting glucose, and waist circumference, which means an honest forward-looking prior is positive drift in the wrong direction. By setting the prior to zero, the cohort does not receive credit for avoiding the slow background deterioration that would otherwise have occurred. Future versions may relax this to a population-baseline prior, with appropriate documentation.

§4. Credibility Weighting (Bühlmann)

The cohort-specific observed delta on each marker is credibility-weighted against the conservative prior using the Bühlmann (1967) credibility formula. This section specifies the formula, the choice of , the derivation of from variance components, and three worked examples spanning the operationally relevant cohort-size range.

§4.1 The Bühlmann formula

The Bühlmann credibility weight is:

> CREDIBILITY WEIGHT
>
> $$Z = n / (n + k)$$
>
> *where n is exposure in person-quarters and k is the credibility constant.*

The credibility-weighted estimate of the cohort's true per-employee improvement on marker is:

> CREDIBILITY-WEIGHTED DELTA
>
> $$\hat{\Delta}_m = Z \cdot \Delta_{m,observed} + (1 - Z) \cdot \Delta_{m,prior}$$
>
> *In v1.0, $\Delta_{prior} = 0$ for all markers, so $\hat{\Delta}_m = Z \cdot \Delta_{m,observed}$.*

The interpretation is the standard one: is the fraction of the cohort-specific signal that is credible given the exposure, with the remaining (1 −) fraction defaulting to the prior. Small cohorts receive less credit for their observed deltas; large cohorts observed over multiple quarters receive more.

§4.2 The Bühlmann constant *k*

In Bühlmann's framework, is the ratio of expected process variance to variance of hypothetical means:

> BÜHLMANN CONSTANT — DEFINITION
>
> $$k = E[Var(X|\Theta)] / Var[E(X|\Theta)]$$
>
> *numerator: within-cohort (process) variance of the per-PQ marker delta. denominator: between-cohort (structural) variance in true cohort-mean drift.*

For v1.0, Metra adopts a unified = **400 person-quarters** across all four markers. The derivation is presented below for the weight marker (the most data-rich case) and is asserted to be a conservative anchor for the remaining three markers, with empirical re-estimation per marker scheduled for v1.1 once production cohort variance is sufficiently observed.

§4.3 Derivation of *k* from variance components

Process variance (within-cohort, per-PQ). The Diabetes Prevention Program (Knowler et al., NEJM 2002) and Look AHEAD (Look AHEAD Research Group, NEJM 2013) report person-level quarterly weight-change standard deviations in the range of 1.6 to 2.0 kg for actively engaged working-age adults. Taking the midpoint and converting to pounds:

PROCESS VARIANCE — WEIGHT, PER PERSON-QUARTER

OBSERVED SD (PER-PQ WEIGHT CHANGE) $\sigma_{PQ} \approx 1.8$ kg = 3.97 lbs ≈ 4.0 lbs

```
Var(X|θ)_PQ ≈ (4.0 lbs)² = 16.0 lbs² per PQ
```

This is the noise floor: even a cohort with zero true drift will produce per-PQ observed deltas with this much spread.

Variance of hypothetical means (between-cohort). The between-cohort variance is the variance of true per-employee mean drift across the population of working-age employer cohorts. NHANES (NCHS, 2017–2020 cycle) population estimates of working-age weight change yield a cohort-on-cohort SD for annual mean drift on the order of 0.3 to 0.5 lbs per year for sufficiently large strata. Converting to a per-PQ quantity:

VARIANCE OF HYPOTHETICAL MEANS — WEIGHT, PER PERSON-QUARTER

BETWEEN-COHORT SD (ANNUAL MEAN DRIFT) $\sigma_{annual} \approx 0.4$ lbs / yr

CONVERT TO PER-PQ $\sigma_{PQ\text{-}mean} = 0.4 / 4 = 0.10$ lbs per PQ

VAR[E(X|Θ)]_PQ

```
Var[E(X|θ)]_PQ ≈ (0.10 lbs)² × 4 ≈ 0.04 lbs² per PQ
```

(The factor of 4 reflects that quarterly mean drifts compound to the annual SD reported by NHANES; the per-PQ structural variance is the annual variance divided by 4 under an i.i.d. quarterly drift assumption, which is itself conservative — autocorrelation in cohort drift would reduce this further, increasing k.)

Ratio.

BÜHLMANN CONSTANT — WEIGHT MARKER

k_{weight} = *16.0 / 0.04* = ***400 person-quarters***

Application to the other three markers. The analogous variance-ratio analysis on systolic BP, fasting glucose, and waist circumference, using DPP and Look AHEAD process-variance estimates and NHANES between-cohort variance, yields values in the range of 350 to 480 PQ. v1.0 adopts = 400 PQ as the unified constant across all four markers. The implication is that for markers where the true is greater than 400 PQ (i.e., the unified value is slightly too low), the v1.0 Forecast may slightly over-credit the cohort signal on those markers. Two mitigations are designed in to absorb this:

1. The per-marker confidence interval (§5) widens with cohort variance and naturally absorbs additional uncertainty.
2. The 12-month forecast applies the lag haircut λ = 0.50 (§6), which provides additional cushion against any over-credit at long horizon.

v1.1 will publish per-marker values estimated directly from accumulated production data, with the four-marker variance decomposition recorded for review.

§4.4 Worked examples across the cohort-size range

The following three examples span the operationally relevant cohort sizes. The 322-employee cohort matches the Metra demonstration corporate population used to produce the published Sample Workforce Exposure Forecast.

EXAMPLE A — SMALL COHORT, PARTIAL OBSERVATION WINDOW

COHORT N = 50 employees, 6 months of valid weight observation

PERSON-QUARTERS n = 50 × 2 = 100 PQ

```
Z = 100 / (100 + 400) = 0.200
```

20% of the observed cohort weight delta is credible signal; 80% defaults to the zero-improvement prior. A 50-employee cohort with a six-month observation window is properly humbled by the credibility math.

EXAMPLE B — MEDIUM COHORT, FULL ANNUAL WINDOW (MATCHES DEMO)

COHORT N = 322 employees, 12 months of valid weight observation

PERSON-QUARTERS n = 322 × 4 = 1,288 PQ

```
Z = 1,288 / (1,288 + 400) = 0.763
```

76.3% of the observed cohort weight delta is credible signal; 23.7% defaults to the prior. This is the credibility weight used in the published Sample Workforce Exposure Forecast.

EXAMPLE C — LARGE COHORT, FULL ANNUAL WINDOW

COHORT N = 1,000 employees, 12 months of valid weight observation

PERSON-QUARTERS n = 1,000 × 4 = 4,000 PQ

```
Z = 4,000 / (4,000 + 400) = 0.909
```

90.9% credibility on the cohort signal. At this exposure level, the prior contributes less than 10% of the credibility-weighted estimate.

The shape of the credibility curve is intentional. A 50-employee cohort with six months of observation is well below the carrier's customary credibility threshold for employer-specific experience pricing; the v1.0 methodology agrees, crediting only 20% of the observed signal. A 1,000-employee cohort with a full year of valid biometric observation reaches 90% credibility on the biometric signal, comparable to the credibility a carrier would assign to a year of fully-credible claims experience on the same population.

§4.5 Use of the credibility weight downstream

The credibility-weighted delta $\hat{\Delta}$ is the quantity propagated forward into the per-marker exposure formula (§5) and the forecast formulas (§6). It is not separately reported on the Forecast cover. The Forecast does report the per-marker credibility weight in the methodology footnote so the reading actuary can independently verify the application.

§5. Per-Marker Exposure Estimation and Confidence Interval Construction

Per-marker exposure is the product of cohort size, credibility-weighted per-employee improvement, and per-unit cost coefficient. Confidence intervals on per-marker and total exposure are constructed via percentile bootstrap on the cohort. This section specifies the point estimate, the sources of uncertainty, the bootstrap procedure, and the worked example that produces the per-marker bands published on the Forecast.

§5.1 Per-marker point estimate

> PER-MARKER EXPOSURE (ANNUALIZED)
>
> $$E_m = N \cdot \hat{\Delta}_m \cdot c_m$$
>
> *with Δ_m the credibility-weighted per-employee improvement from §4, and c_m the per-unit cost coefficient from §9.*

Total exposure is the sum across the four markers:

> TOTAL EXPOSURE (ANNUALIZED)
>
> $$E_{total} = \Sigma_m E_m$$

§5.2 Sources of uncertainty

Three sources of uncertainty enter the per-marker exposure estimate:

1. **Sampling error on $\Delta_{m,\text{observed}}$**: the cohort-mean observed delta is a sample mean and has its own sampling distribution. The standard error scales as $\sigma_\Delta / \sqrt{N}$. This is the dominant source of CI width at small cohort sizes.
2. **Coefficient uncertainty on c_m**: each published coefficient has a confidence band in the source literature; v1.0 anchors each c_m at the lower bound of the published range (see §9). The point estimate is therefore already conservative. v1.0 propagates the coefficient as a point value and discloses the source-literature range in §9; v1.1 will propagate the coefficient as a uniform distribution over the disclosed range.
3. **Credibility weight Z**: deterministic given N and k, so it does not contribute to the CI width.

§5.3 Confidence interval procedure — percentile bootstrap

The 95% confidence interval on per-marker and total exposure is constructed via the percentile bootstrap. The procedure is specified below.

> BOOTSTRAP CI PROCEDURE
>
> *For b = 1 to B (with B = 10,000):*
>
> 1. *Resample N employees from the cohort with replacement, recording indices* I^b.
> 2. *For each marker m, compute* $\Delta_{m,observed}^{(b)}$ *from the per-employee deltas restricted to* I^b.
> 3. *Apply credibility weighting:* $\hat{\Delta}_m^{(b)} = Z \cdot \Delta_{m,observed}^{(b)}$.
> 4. *Compute per-marker exposure* $E_m^{(b)} = N \cdot \hat{\Delta}_m^{(b)} \cdot c_m$.
> 5. *Compute total exposure* $E_{total}^{(b)} = \Sigma_m E_m^{(b)}$.
>
> *The 95% CI on each quantity is the [2.5%, 97.5%] percentile of its bootstrap distribution.*

Why this procedure. Three reasons.

1. **No parametric assumption on the per-employee delta distribution.** Workforce biometric improvement distributions are typically right-skewed (a tail of strong improvers, a mass near zero, occasional regressors). A normal-distribution CI would systematically under-cover. The bootstrap inherits the empirical distribution.
2. **Cohort-size sensitivity is automatic.** Bootstrap variance scales as $1/\sqrt{N}$; the CI naturally widens for small cohorts without any additional parameter.
3. **Cross-marker correlation is preserved without assumption.** Step 1 resamples the *same* employee indices across all markers within each bootstrap iteration. The empirical correlation between markers (high-improver employees showing up coherently across markers) is preserved. The total-exposure CI does not require a correlation matrix to be assumed.

Choice of B = 10,000. Standard percentile-bootstrap practice (Efron & Tibshirani, 1993) recommends B in the 1,000–10,000 range for 95% CIs. v1.0 uses 10,000 to ensure the [2.5%, 97.5%] endpoints are estimated with Monte Carlo error well below the underlying sampling error, even for small cohorts where the underlying CI is wide. Runtime cost is negligible for cohorts up to several thousand.

§5.4 Worked example — 322-employee cohort, weight marker

Using the demonstration corporate population from the Sample Workforce Exposure Forecast:

WORKED EXAMPLE — WEIGHT MARKER, DEMO COHORT

COHORT N = 322, 12 months observation, full per-quarter coverage

OBSERVED DELTA $\Delta_{weight,obs}$ = 5.2 lbs per-employee average reduction

CREDIBILITY WEIGHT Z = 1,288 / (1,288 + 400) = 0.763

CREDIBILITY-WEIGHTED DELTA $\tilde{\Delta}_{weight}$ = 0.763 × 5.2 = 3.97 lbs

COEFFICIENT c_{weight} = $68 per lb per person-year

POINT ESTIMATE

```
E_weight = 322 × 3.97 × $68 = $86,939 / yr
```

95% CI FROM B = 10,000 BOOTSTRAP ITERATIONS

```
95% CI: [$72,400 – $102,180]
```

The point estimate is the credibility-weighted, conservative-anchor exposure on the weight marker alone for the demonstration cohort. The CI width reflects the sampling variability on the cohort-mean delta; it does not include coefficient uncertainty (which is held at the published lower bound — see §9).

The same procedure applied across all four markers and summed (with cross-marker correlation preserved via shared bootstrap indices) yields the total per-marker breakdown and total exposure published on the Forecast cover.

§6. Forecast Horizon Construction

The Forecast reports exposure projections at four horizons: 30, 60, and 90 days, and 12 months. The 30/60/90-day projections are produced by linear extrapolation of the observed cohort trajectory with horizon-widening prediction intervals. The 12-month projection adds two additional adjustments — a lag haircut for partial flow-through of biometric improvement to claim spend, and an attrition adjustment reflecting the effective shrinkage of the at-risk cohort over the year. This section specifies all three constructions and defends the choice of horizons against the obvious alternatives.

§6.1 Why these horizons

The four horizons are chosen to align with two operationally distinct uses of the Forecast.

- **30 / 60 / 90 days — current trajectory.** Workforce biometric drift is measurable at quarterly resolution with adequate signal-to-noise (§3). The 30/60/90-day window is the "what is the cohort doing right now" horizon, used by HR and the wellness program operationally and by the broker in mid-cycle conversations.
- **12 months — renewal cycle.** Commercial group medical policies in the U.S. employer-sponsored market overwhelmingly renew annually. The 12-month forecast is the horizon at which the CFO, broker, and actuary make planning and pricing decisions. It is the horizon at which the per-marker dollar figures on the Forecast cover are most directly actionable.

Why no 6-month or 9-month forecast. The biometric-to-claim lag literature (§8) supports a multi-month-to-multi-year lag from biometric improvement to realized claim impact, but it does not temporally resolve a 6-month or 9-month claim-impact estimate to publishable precision. No single published study identifies a defensible 6-month or 9-month coefficient. Producing a forecast at those horizons would require either an interpolation that overstates the precision available in the source literature, or a coefficient assumption that we cannot defend at this document's bar. v1.0 declines the temptation. The 30/60/90/12-month set is the temporal grid the source literature actually supports.

Why no 18-month or 24-month forecast. Two reasons. First, the prediction interval on linear extrapolation widens quadratically with horizon (§6.2 below). At 18 months, the PI band on most realistic cohort trajectories spans both the no-improvement and substantial-improvement scenarios — the forecast loses its ability to differentiate. Second, the linear-extrapolation model assumes the absence of intervening interventions and stable cohort composition. Both assumptions degrade past 12 months. The Society of Actuaries' general guidance on experience-period projection for similar instruments caps the projection horizon at 12–18 months for these same reasons. v1.0 caps at 12.

§6.2 Short-horizon (30/60/90) forecast and prediction interval

For each marker , the observed cohort-mean trajectory is the sequence of quarterly cohort means Δ (t_1), Δ (t_2), ..., Δ (t) where the are the quarter-end observation times. A simple linear regression is fit:

TRAJECTORY LINEAR FIT

$$\Delta_m(t) = \alpha_m + \beta_m \cdot t + \varepsilon$$

with t in days from observation start and residual variance σ_ε^2.

The point projection at horizon is $\hat{\Delta}$ ( ) = $\hat{\alpha}$ + $\hat{\beta}$ · . The prediction interval on the projected cohort mean uses the standard simple-linear-regression PI formula:

PREDICTION INTERVAL — PROJECTED COHORT MEAN AT HORIZON H

$$SE(\hat{\Delta}_m(h)) = \sigma_\varepsilon \cdot \sqrt{[1/k_{obs} + (h - \bar{t})^2 / \Sigma(t_i - \bar{t})^2]}$$

with k_{obs} the number of quarterly observations in the trajectory fit and $\bar{t}$ the mean observation time.

The second term inside the square root grows quadratically as moves away from the centroid of the observation window. At = 30 days the PI is tight; at = 90 days it has visibly widened; at = 365 days the widening is meaningful and is one of two reasons the 12-month forecast carries additional adjustments (§6.3).

The 95% PI on the projected per-marker exposure at horizon is constructed by combining the trajectory PI with the §5 bootstrap procedure: for each bootstrap iteration, the trajectory is re-fit on the resampled cohort, the horizon- projection is produced, and the per-marker exposure at is computed. The [2.5%, 97.5%] percentile of the bootstrap distribution is the reported 95% PI.

§6.3 Twelve-month forecast — additional adjustments

The 12-month forecast applies two adjustments beyond the linear projection: a lag haircut and an attrition adjustment.

§6.3.1 Lag haircut λ

Biometric improvement does not translate to claim-spend reduction instantaneously. The DPP economic analysis (Knowler et al., 2002), the Million Hearts cost-effectiveness modeling (CDC, 2014–2017), and the AHRQ HCUP cost-of-illness analyses all report claim-impact translation timelines that begin to manifest at 6–18 months from sustained biometric change and reach the per-marker actuarial cost-of-illness coefficient over 24–36 months. The v1.0 methodology applies a lag haircut = 0.50 to the 12-month forecast: only half of the projected 12-month biometric improvement is credited against 12-month claim spend.

Defense of = 0.50. Three anchors.

1. The DPP economic analysis reports the per-mg/dL fasting-glucose reduction translates to cost-of-care reduction with a first-year realization fraction near 0.45 of the sustained-state coefficient.
2. The Million Hearts modeling for blood pressure reports a first-year realization fraction near 0.50 for moderate sustained SBP reduction.
3. Society of Actuaries general guidance on novel risk-mitigation instruments recommends discounting first-year projected savings by 30–50% to absorb realization-timing uncertainty.

Taking the median of these anchors yields ≈ 0.50. v1.0 adopts that figure as the unified lag haircut across all four markers, with per-marker decomposition scheduled for v1.1 once accumulated production claim-experience linkage is sufficient to re-estimate each directly.

§6.3.2 Attrition adjustment α

Workforce attrition reduces the effective at-risk cohort over the year. Employees who leave the employer mid-year do not contribute their full annualized exposure delta to the realized 12-month claim spend (their share of the realized impact flows to whichever payer covers them after their departure). v1.0 applies an annualized attrition rate = 0.15 (15% per year), treating attrition as uniform across the year so that the mean exposure-weighted cohort over the year is:

EFFECTIVE ANNUALIZED COHORT

$$N_{eff} = N \cdot (1 - \alpha/2) = N \cdot 0.925$$

under a uniform-attrition assumption, the average employee is enrolled for 50% of the post-attrition exposure window; the annualized integrated cohort exposure is $N \cdot (1 - \alpha/2)$.

Defense of = 0.15. The Bureau of Labor Statistics Job Openings and Labor Turnover Survey (JOLTS) reports total annualized separation rates for U.S. private-sector employers in the 12–18% range across the 2020–2025 period (with cyclical variation). v1.0 adopts the midpoint of this range as the default. The methodology supports per-client override of to a client-specific JOLTS-comparable figure where the employer reports actual workforce turnover — this override is disclosed in the Forecast footnote when applied.

§6.3.3 Twelve-month forecast formula

> TWELVE-MONTH WORKFORCE EXPOSURE FORECAST
>
> $$E_{total}(12mo) = N_{eff} \cdot \Sigma_m [\, \Delta_m(365) \cdot c_m \cdot \lambda \,]$$
>
> with $\Delta_m(365)$ the credibility-weighted, trajectory-projected 12-month per-employee delta on marker m; $N_{eff} = N \cdot 0.925$; $\lambda = 0.50$.

The 95% PI on the 12-month total exposure is constructed via the bootstrap procedure described in §5.3 and §6.2, with the lag haircut and attrition adjustment applied within each bootstrap iteration. The PI is reported on the Forecast cover alongside the point estimate.

§6.4 Worked example — 322-employee cohort, 12-month forecast

WORKED EXAMPLE — 12-MONTH TOTAL FORECAST, DEMO COHORT

COHORT N = 322, 12-mo observation, full quarterly coverage on all four markers

CREDIBILITY WEIGHT Z = 0.763 (from §4 example B)

EFFECTIVE ANNUALIZED COHORT N_eff = 322 × 0.925 = 297.85

PROJECTED PER-EMPLOYEE DELTAS AT H = 365 DAYS (CREDIBILITY-WEIGHTED)

```
Δ̂weight(365) = 0.763 × 6.1 lbs = 4.65 lbs Δ̂BP(365) = 0.763 × 5.8 mmHg = 4.43 mmHg
Δ̂BG(365) = 0.763 × 7.4 mg/dL = 5.65 mg/dL Δ̂waist(365) = 0.763 × 4.3 cm = 3.28 cm
```

PER-MARKER 12-MONTH EXPOSURE WITH LAG HAIRCUT Λ = 0.50

```
E_weight(12mo) = 297.85 × 4.65 × $68 × 0.50 = $47,083 E_BP(12mo) = 297.85 × 4.43 × $150
× 0.50 = $98,953 E_BG(12mo) = 297.85 × 5.65 × $50 × 0.50 = $42,066 E_waist(12mo) =
297.85 × 3.28 × $30 × 0.50 = $14,650
```

TOTAL 12-MONTH EXPOSURE (POINT ESTIMATE)

```
E_total(12mo) = $202,752 / yr
```

95% PI FROM BOOTSTRAP (B = 10,000)

```
95% PI: [ $164,400 – $245,100 ]
```

The PI band is approximately ±20% of the point estimate at 12 months — wider than the per-marker bands at 30/60/90 days, reflecting the additional trajectory-projection and lag uncertainty at the longer horizon. This is the figure carried into the renewal conversation as the Forward 12-Month Healthcare Exposure Avoided.

The three preceding sections — §4, §5, §6 — constitute the mathematical core of the methodology. Every quantity reported on a generated Workforce Exposure Forecast is derivable from these sections combined with the per-marker coefficients in §9.

§7. METRA Score Composition

The METRA Score (0–100) is the per-employee summary metric used in the platform's user-facing surfaces. It does not enter the workforce exposure forecast directly — the forecast uses raw biometric deltas, not the score — but is documented here for completeness because it is referenced on the Forecast cover as a cohort-distribution summary.

METRA SCORE — PER-EMPLOYEE, 0–100

$$Score = 0.30 \cdot S_{trend} + 0.25 \cdot S_{BP} + 0.25 \cdot S_{body} + 0.20 \cdot S_{consistency}$$

each component is a 0–100 sub-score; the weights sum to 1.00.

- $\mathbf{S_{trend}}$ (weight 0.30): composite of recent weight trajectory and metabolic-age delta over the last 90 days.
- $\mathbf{S_{BP}}$ (weight 0.25): systolic and diastolic banded against AHA categories with bonus for sustained improvement.
- $\mathbf{S_{body}}$ (weight 0.25): waist circumference, BMI band, and body-composition deltas.
- $\mathbf{S_{consistency}}$ (weight 0.20): cadence of valid biometric and food-log entries over the last 30 days.

The score is computed nightly per active enrolled employee. Cohort-level statistics (mean score, score distribution by risk band) are aggregated under the de-identification protocol described in §3 and the Forecast template (§10).

§8. Biometric-to-Claim Pathway and Lag Literature

The exposure forecast rests on the empirical claim that sustained biometric improvement at the workforce level causally precedes, by a measurable lag window, a reduction in healthcare claim spend on the affected population. This section summarizes the source literature for that claim and identifies where the v1.0 lag haircut = 0.50 (§6.3.1) is anchored.

Glucose pathway. Tabák et al. (Lancet, 2009) documented multi-year fasting-glucose and insulin-sensitivity trajectories preceding type 2 diabetes diagnosis, with measurable upstream deterioration observable 5–10 years before the diagnostic event. Knowler et al. (NEJM, 2002 — Diabetes Prevention Program) quantified the modifiability of that trajectory and the per-unit cost-of-care implication. ADA's (2017 update) provides the population-level annualized cost figures from which the per-mg/dL coefficient (§9) is derived. The first-year realization fraction of the per-mg/dL improvement on claim spend is approximately 0.45 in the DPP cost-effectiveness modeling.

Blood pressure pathway. Moran et al. (Annals of Internal Medicine, 2015) modeled per-mmHg cost-effectiveness of sustained systolic-BP reduction under the 2014 hypertension guidelines, with first-year realization fraction near 0.50. The CDC Million Hearts initiative provides the implementation-context anchoring for the per-mmHg coefficient. The first-year-realized-savings fraction sits at the upper end of the 0.45–0.55 range.

Body composition pathway. Finkelstein et al. (Health Affairs, 2009; 2014 update) and Tsai et al. (Obesity Reviews, 2011) provide the per-pound cost-of-obesity annualized figures. AHRQ HCUP statistical briefs and Boudreau et al. (Population Health Management, 2009) provide the per-cm waist-circumference cost-of-illness anchoring. First-year realization fractions for sustained weight and waist reduction are reported in the 0.40–0.55 range across these sources.

Application to . The four-marker first-year realization fractions cluster in the 0.40–0.55 range. v1.0 adopts the rounded median = 0.50 as the unified haircut, consistent with the SOA guidance band for novel risk-mitigation instruments. v1.1 will publish per-marker derived from accumulated production claim-experience linkage. The unified value in v1.0 is conservative against the upper-end markers (BP at 0.55) and at-anchor against the lower-end markers (glucose at 0.45).

The biometric-to-claim pathway literature is established. What is new in the v1.0 methodology is not the existence of the pathway — it is the application of an actuarial credibility framework to continuously observed workforce biometric signal, in a buyer-side instrument that produces a defensible per-client forward exposure figure.

§9. Per-Marker Cost Coefficients

The four per-unit annual cost coefficients used in the per-marker exposure formula (§5) are reproduced below. Each coefficient is anchored at the conservative (lower-bound) end of the published source-literature range. The v1.0 methodology treats each coefficient as a point estimate at its lower-bound anchor; v1.1 will propagate the published range as a uniform distribution within the bootstrap.

MARKER	UNITS	COEFFICIENT (USD / PERSON-YR / UNIT)	PUBLISHED RANGE	PRIMARY SOURCE
Weight	per lb	**$68**	$68 – $145	Finkelstein et al., *Health Affairs* (2009; 2014 update); Tsai et al., *Obesity Reviews* (2011).
Systolic BP	per mmHg	**$150**	$100 – $300	Moran et al., *Annals of Internal Medicine* (2015); CDC Million Hearts (2014–2017).
Fasting glucose	per mg/dL	**$50**	$40 – $100	Knowler et al., *NEJM* (2002) — DPP; ADA *Economic Costs of Diabetes in the U.S.* (2017).
Waist	per cm	**$30**	$20 – $60	AHRQ HCUP Statistical Briefs; Boudreau et al., *Population Health Management* (2009).

The choice to anchor at lower bound is intentional. It means that the published Forecast point estimate is the lowest defensible figure in the literature; any reviewing actuary who substitutes the midpoint of the published range for any coefficient will arrive at a higher figure than the Forecast reports. This is the doctrine: the per-client Forecast is intended to under-report the upper bound of plausible exposure avoidance, not over-report it.

Coefficient version pinning. Each coefficient in the table is pinned to the version of this methodology in which it appears. A coefficient revision constitutes a methodology version bump (v1.0 → v1.1 → ...) and triggers re-attestation. The dashboard's running Forecast generator inserts the active methodology version in the Forecast footer; an actuary reading a generated Forecast can always look up which coefficient version was applied.

§10. Aggregation to Workforce Total

The full aggregation formula combining §3 through §9 is reproduced here for completeness. For a single Forecast generation pass on a cohort of size with credibility weight , projected at horizon :

WORKFORCE EXPOSURE — SHORT HORIZON (H ∈ {30, 60, 90} DAYS)

$$E_{total}(h) = N \cdot \Sigma_m [Z \cdot \hat{\Delta}_m(h) \cdot c_m]$$

with $\hat{\Delta}_m(h) = \hat{\alpha}_m + \hat{\beta}_m \cdot h$ from the trajectory fit.

WORKFORCE EXPOSURE — 12-MONTH RENEWAL HORIZON

$$E_{total}(365) = N \cdot (1 - \alpha/2) \cdot \Sigma_m [Z \cdot \hat{\Delta}_m(365) \cdot c_m \cdot \lambda]$$

95% confidence intervals on each are produced by the bootstrap procedure of §5.3, with the trajectory fit re-performed on each resampled cohort. The Forecast cover reports the point estimate at the upper-right and the CI band immediately below, with per-marker decomposition in the body of the Forecast.

Snapshot integrity. Every generated Forecast embeds in its footer a snapshot digest comprising: (a) the methodology version (v1.0), (b) the values of , , , , (c) the four coefficient values, (d) the cohort size and credibility weight , (e) the date range of the observation window, and (f) a hash of the de-identified per-employee delta vectors used in the bootstrap. The digest allows any reviewing actuary to confirm that two Forecasts purporting to describe the same cohort do in fact rest on the same underlying data.

§11. Limitations and Known Constraints

v1.0 of this methodology is published with the following limitations openly named. Each is a defensible v1.0 choice; each is also a candidate for v1.1 refinement.

L1 **L1. Unified *k* across markers.**

The Bühlmann constant = 400 PQ is applied uniformly across all four markers. The variance-component derivation (§4.3) is rigorous for weight; for the other three markers the analogous derivation yields values in the range 350–480 PQ. Adopting a unified value may slightly over-credit cohort signal on the higher- markers. The bootstrap CI (§5) and the lag haircut (§6) absorb this conservatism. v1.1 will publish per-marker .

L2 **L2. Zero prior on biometric drift.**

The credibility-weighted estimate defaults the (1 −) fraction to a zero-improvement prior, rather than to the small positive deterioration NHANES reports for working-age U.S. adults. This is intentionally conservative: the cohort does not receive credit for avoiding background deterioration. v1.1 may move to a population-baseline prior, with documented effect on the point estimate.

L3 **L3. Coefficient point estimate.**

Each is applied as a point value at the published lower bound. v1.0 does not propagate coefficient uncertainty into the bootstrap. The reported CI therefore reflects sampling uncertainty on the cohort delta only, not coefficient uncertainty. v1.1 will propagate each as a uniform distribution over its published range.

L4 **L4. Linear trajectory model.**

The short-horizon and 12-month forecasts use a simple linear regression on quarterly cohort means. Real cohort biometric trajectories exhibit non-linearities (rapid initial improvement, asymptotic flattening, plateau effects). The linear model is conservative at long horizons (it generally projects continued improvement rather than asymptote), which is a relevant caveat for the 12-month forecast. v1.1 may move to a piecewise or saturating-curve model.

L5 **L5. Unified lag haircut.**

The 12-month lag haircut = 0.50 is unified across markers. Per-marker realization fractions in the source literature span 0.40–0.55. v1.1 will publish per-marker .

L6 L6. Uniform attrition assumption.

The attrition adjustment = 0.15 is applied uniformly across the year. Real workforce attrition is non-uniform (seasonal, role-dependent, tenure-dependent). The uniform-attrition formula $_{eff}$ = · (1 − /2) is a first-order approximation. v1.0 allows per-client override of ; v1.1 may support a per-client attrition curve.

L7 L7. No carrier-experience integration.

v1.0 produces a stand-alone biometric-signal forecast. It does not combine the biometric signal with the carrier's incurred-claim experience credibility for the same cohort. Integration of the two — the buyer-side instrument folded into the carrier-side credibility math — is the explicit invitation in the attestation protocol (§12) but is not specified by this methodology document.

L8 L8. Cross-marker correlation via shared bootstrap indices.

The total-exposure CI is constructed by resampling the same employee indices across all markers, which preserves empirical correlation. This is the cleanest non-parametric option but does not separately publish a per-pair correlation matrix. A reviewing actuary who wants to inspect the implied correlation structure can request the per-iteration bootstrap output through the same channel that requests the cohort delta vectors.

L9 L9. Cohort floor below carrier-side credibility threshold.

The cohort floor of N ≥ 15 (company-wide) and N ≥ 5 (scoped) is set at the de-identification floor for HR-facing dashboards, not at the carrier-side experience-credibility floor (which is typically much higher). The Forecast generated at small cohort sizes will carry wide CIs and low credibility weights, and the methodology surfaces both — but the small-cohort Forecast should not be read as a substitute for a credibility-weighted carrier rate.

---

§12. Attestation Protocol

What attestation is. A credentialed actuary's signed statement that they have read this methodology document at the published version, applied their professional judgment to the formulas, derivations, coefficients, and limitations, and find the document defensible at the publishing version. The attestation is recorded at usemetra.com/methodology/attestations/ with the actuary's name, credential, firm, the version attested, and the attestation date.

What attestation is not. An attestation does not constitute a regulatory filing, a rate certification, an actuarial opinion under ASOP 41 or any successor, or a representation about any specific employer's exposure. It is a professional statement on the methodology document itself, not on any particular Forecast produced from it.

Tier model. Each generated Workforce Exposure Forecast carries one of three signature blocks on the cover, conditional on the state of the employer's attestation roster at generation time:

- **Tier 1 — Employer-retained actuary.** Where the employer has an actuary of record (typically through their broker firm or as an independent retained consultant) who has attested to the current methodology version, that actuary's signature block is on the cover. The Forecast is signed by the credentialed professional engaged by the buyer.
- **Tier 2 — Broker-firm actuary.** Where the broker firm has a credentialed actuary on staff who has attested to the current version, and the employer has not designated an alternate signer, the broker-firm actuary's signature block is on the cover. The Forecast is signed by the credentialed professional engaged by the buyer's broker.
- **Tier 3 — Methodology version attestation only.** Where no employer-side or broker-side actuary is currently attesting the methodology version, the cover carries no individual signer. Instead, the cover displays the list of currently-active attestations on the methodology version itself (with names, credentials, firms, attestation dates), and explicitly labels the Forecast as *unsigned at the per-client level*. The Metra platform separately offers, at no additional cost to the employer, engagement with a vetted independent actuary panel to convert a Tier 3 Forecast to a Tier 1 signed Forecast. Eligibility for this credit is described in the platform's billing terms; it is not specified by this methodology document.

Re-attestation on version bump. Every change to this methodology document that affects any formula, coefficient, or constant constitutes a version bump (v1.0 → v1.1 → ...). All attestations are version-pinned. An attester whose attestation is on a superseded version remains listed in the historical attestation roster for that version, but is not listed as current until they review and re-attest the new version. The platform notifies all version-pinned attesters when a new version publishes.

How to attest. Credentialed actuaries (ASA, FSA, MAAA, FCAS, ACAS) may submit an attestation by reading the published version of this document and contacting Metra at care@usemetra.com with the subject line “Methodology Attestation — v1.0” (substituting current version). The attestation submission carries the actuary's name, credential, firm, credential number, an optional firm letterhead document, and a statement that they have reviewed the methodology at the named version and find it defensible. Compensation arrangements (where any exist) are disclosed on the public attestation roster.

§13. Version History and Change Control

Every published version of this methodology is preserved at a permanent URL of the form `usemetra.com/methodology/v /`. The current version's content is mirrored at `usemetra.com/methodology/` as a convenience for the most recent reader. Superseded versions remain readable at their version URLs so that attestations against historical versions remain auditable.

VERSION	PUBLISHED	STATUS	SUMMARY OF CHANGES
v1.0	2026	Current — open for review & attestation	Initial publication. Cohort construction, Bühlmann credibility (k = 400 PQ unified), per-marker exposure with bootstrap CI (B = 10,000), short-horizon trajectory PI, 12-month forecast with $\lambda = 0.50$ lag haircut and $\alpha = 0.15$ attrition adjustment, four-marker coefficient table at literature lower bound.
v1.1	planned	Pending commissioned external review	Per-marker k values from production data; per-marker λ_m from claim-experience linkage; coefficient uncertainty propagated as uniform distributions within bootstrap; piecewise or saturating trajectory model; population-baseline prior on biometric drift; commissioned external actuarial review and resulting refinements.

§14. BIBLIOGRAPHY

1. **Bühlmann, H.** (1967). Experience rating and credibility. *ASTIN Bulletin*, 4(3), 199–207. — Foundational credibility-theory paper from which the $Z = n/(n+k)$ formula is taken.
2. **Bühlmann, H., & Gisler, A.** (2005). *A Course in Credibility Theory and its Applications*. Springer. — Standard graduate reference for the credibility framework applied in §4.
3. **Klugman, S. A., Panjer, H. H., & Willmot, G. E.** (2019). *Loss Models: From Data to Decisions* (5th ed.). Wiley. — SOA-curriculum reference for credibility, simulation, and bootstrap methods relevant to §4 and §5.
4. **Efron, B., & Tibshirani, R. J.** (1993). *An Introduction to the Bootstrap*. Chapman & Hall. — Reference for the percentile bootstrap procedure used in §5.3.
5. **Society of Actuaries.** (Various years). Group & Health Issue Briefs; credibility guidance for novel risk-mitigation instruments. — Reference for the 30–50% first-year realization discount used in the derivation of λ in §6.3.1.
6. **Knowler, W. C., et al.** (2002). Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin. *New England Journal of Medicine*, 346(6), 393–403. — Diabetes Prevention Program; per-mg/dL realization fraction cited in §6.3.1 and §8.

7. **Tabák, A. G., et al.** (2009). Trajectories of glycaemia, insulin sensitivity, and insulin secretion before diagnosis of type 2 diabetes: an analysis from the Whitehall II study. *The Lancet*, 373(9682), 2215–2221. — Upstream biometric-signal evidence cited in §8.
8. **Look AHEAD Research Group.** (2013). Cardiovascular effects of intensive lifestyle intervention in type 2 diabetes. *New England Journal of Medicine*, 369(2), 145–154. — Within-cohort weight-change variance estimates used in the derivation of k in §4.3.
9. **Moran, A. E., et al.** (2015). Cost-effectiveness of hypertension therapy according to 2014 guidelines. *Annals of Internal Medicine*, 162(10), 661–671. — Per-mmHg coefficient anchor (§9) and first-year realization fraction (§6.3.1).
10. **Centers for Disease Control and Prevention.** (2014–2017). *Million Hearts Initiative — Cost-Effectiveness Modeling*. — Per-mmHg implementation-context anchor cited in §9.
11. **American Diabetes Association.** (2018). Economic costs of diabetes in the U.S. in 2017. *Diabetes Care*, 41(5), 917–928. — Per-mg/dL coefficient anchor cited in §9.
12. **Finkelstein, E. A., et al.** (2009). Annual medical spending attributable to obesity: payer- and service-specific estimates. *Health Affairs*, 28(5), w822–w831 (with 2014 update). — Per-lb coefficient anchor cited in §9.
13. **Tsai, A. G., et al.** (2011). Direct medical cost of overweight and obesity in the USA: a quantitative systematic review. *Obesity Reviews*, 12(1), 50–61. — Per-lb coefficient range cited in §9.
14. **Agency for Healthcare Research and Quality.** (Various years). *HCUP Statistical Briefs on metabolic syndrome cost*. — Per-cm waist coefficient anchor cited in §9.
15. **Boudreau, D. M., et al.** (2009). Health care utilization and costs by metabolic syndrome risk factors. *Population Health Management*, 12(6), 317–324. — Per-cm waist coefficient range cited in §9.
16. **National Center for Health Statistics.** (2017–2020 cycle). *National Health and Nutrition Examination Survey (NHANES)*. — Between-cohort variance estimates for working-age weight, BP, glucose, and waist trajectory used in the derivation of k in §4.3 and the prior discussion in §3 and §11.
17. **U.S. Bureau of Labor Statistics.** (2020–2025). *Job Openings and Labor Turnover Survey (JOLTS)*. — Annualized separation-rate range cited in the derivation of α in §6.3.2.