METRA · WORKFORCE EXPOSURE FORECAST

Forward Healthcare Exposure for Acme Corp Demo

SAMPLE · DEMO COHORT

Report date: May 28, 2026 Methodology v1.0 Cohort N = 325

This document is a credibility-weighted forward forecast of healthcare exposure avoided across the Acme Corp Demo workforce, computed from sustained biometric improvement on four metabolic markers (weight, systolic blood pressure, fasting glucose, waist circumference) using the lower-bound per-marker cost coefficients published in the Metra Methodology. The CFO can plan against it. The CHRO can govern with it. The broker can carry it into the carrier conversation. The actuary can fold it into their credibility math.

12-MONTH FORWARD EXPOSURE AVOIDED · 95% CI

$121,864

95% CI: $96,344 — $147,461 · 10,000-draw percentile bootstrap on the per-employee monetary-delta vector, employee-level resampling preserving cross-marker correlation

DEMO SIGNER — ILLUSTRATIVE ENTRY

Marcus Chen, FSA, MAAA DEMO

Chen Actuarial Consulting

Demonstration signing actuary used on the published Sample Workforce Exposure Forecast. No live engagement. Methodology v1.0 is the same document any live signing actuary would attest to.

Sample document — signing engagement pending for live forecasts

Cohort & Observation Window

SAMPLE · DEMO COHORT

Enrolled employees (cohort N)	325
Employees with ≥1 marker delta (contributing)	323
Median observation window (months)	12.0
Attrition assumption (12-month)	15.0%
Effective forward-looking N (N · [1 − α])	276.25

Credibility Weighting (Bühlmann)

Cohort credibility weight applied to all forward exposure figures:

```
Z = N / (N + k)  =  325 / (325 + 400)  =  0.4483
```

The Bühlmann credibility constant **k = 400** is the disclosed value in Metra Methodology v1.0, §3.1. It is calibrated against the carrier-side variance-of-the-hypothetical-means / expected-process-variance ratio for population-level metabolic exposure in mid-sized commercial books. The reported exposure is the credibility-weighted complement: Z × observed + (1 − Z) × the manual prior, which in v1.0 is the no-improvement baseline (zero), so all forecast totals are scaled by Z.

Forward Forecast — 30 / 60 / 90 day & 12-month

HORIZON	POINT ESTIMATE (Z-WEIGHTED)	95% CI LOW	95% CI HIGH
30 days	**$10,155**	$8,029	$12,288
60 days	**$20,311**	$16,057	$24,577
90 days	**$30,466**	$24,086	$36,865
12 months	**$121,864**	$96,344	$147,461

Horizon construction: $\mathbf{E(t) = N_{eff} \cdot Z \cdot \hat{\mu} \cdot (t / 12)}$, where $\hat{\mu}$ is the per-employee annualized monetary delta. The biometric-to-claim translation lag is held inside the per-marker coefficient calibration — see Methodology §3.3 (Forecast Horizon Construction) and §5 (Biometric-to-Claim Pathway).

Per-Marker Composition (12-month, Credibility-Weighted)

MARKER	AVG. IMPROVEMENT	$/UNIT	12-MO EXPOSURE
Weight	−2.9 lbs	$68	**$24,999**

Systolic BP	−3.9 mmHg	$150	**$73,493**
Fasting Glucose	−2.8 mg/dL	$50	**$17,406**
Waist Circumference	−1.6 cm	$30	**$5,965**
Total	—	—	**$121,863**

Each $/unit coefficient is the lower-bound defensible value from published actuarial and government economics literature; see Methodology §4 (Per-Marker Cost Coefficients) and the Actuarial Brief multi-marker appendix at usemetra.com/actuarial-brief.pdf.

Confidence Interval & Reproducibility

Bootstrap protocol. The 95% CI is constructed by drawing **B = 10,000** employee-level resamples (with replacement) from the contributing per-employee monetary-delta vector, computing the resample mean, then taking the 2.5th and 97.5th percentiles of the resample distribution. Because each draw selects a whole employee row, cross-marker correlation is preserved — no independence assumption is imposed between weight, BP, glucose, and waist improvement. The reported CI is the horizon-scaled, credibility-weighted, attrition-adjusted band, matching the headline metric.

Reproducibility. The bootstrap PRNG (Mulberry32) is seeded from the SHA-256 digest of the sorted input vector. Re-running this report against the same cohort snapshot yields bit-identical CI bounds, which a reviewing carrier-side actuary can verify.

Limitations

- The 12-month forward figure is a forecast of *healthcare exposure avoided* against the no-improvement counterfactual — it is not a guarantee of paid-claims savings in the same period, because the biometric-to-claim translation lag (Methodology §5) places a fraction of avoided exposure in subsequent plan years.
- The cohort credibility weight Z = N/(N+k) collapses to zero at N = 0 and approaches unity at N ≫ k. Reports issued on cohorts below the 15-employee floor (§6) are not published at all; the dashboard returns a cohort-floor suppression instead.
- The per-marker coefficients are calibrated to the U.S. commercially-insured employed population. Forecasts on Medicare-eligible, Medicaid, or expat populations require coefficient recalibration before the report is signable.
- The attrition assumption α = 15% is a methodology default; client-specific historical attrition can be substituted in the company forecast settings, which will scale all horizon totals proportionally.

Computed against **Metra Methodology v1.0** (usemetra.com/methodology/v1.0/). Reading guide: usemetra.com/methodology-companion.pdf · Actuarial brief: usemetra.com/actuarial-brief.pdf
Snapshot digest (SHA-256): `3e79a01a2724838351fe39006729a7a06aa6e45686051837bf0385cacdf304d1`