PREPARED FOR ACTUARIAL REVIEW · COMPLETE EDITION

The Metra Actuarial Brief.

Metra is the employer-side forecasting instrument for workforce healthcare exposure — the forward analytic capability the carrier already has and the buyer side has been missing. This is the complete brief: the whole instrument in one document, from how Metra measures workforce metabolic movement, to how it calibrates that movement against the group's own paid claims, to how the result is attested.

Metra Healthcare Intelligence · Actuarial Brief — Complete Edition · Methodology v1.0 + v2.1, unified ·

Published with full assumption disclosure

This document gathers the entire Metra forecasting instrument into a single brief for actuarial review. Until now the methodology has been published one version at a time — the Workforce Exposure Forecast under Methodology v1.0, and the Workforce Claims Forecast / Experience Calibration under v2.0 and then v2.1 — each shipped as its own version-pinned brief. A reviewer who wanted the whole picture had to read three documents and assemble the connective tissue themselves. The complete edition supplies that connective tissue, presenting each methodology version as a clearly-labeled part while preserving the version pinning that keeps every attestation honest.

The instrument has two analytic surfaces and one attestation discipline, organized here into four parts. **Part I** establishes how Metra measures workforce metabolic movement and projects it forward as healthcare exposure. **Part II** sets out the multi-marker cost coefficients that turn that movement into dollars. **Part III** runs the renewal math buyer-side — credibility-blends the group's own paid experience against the carrier manual rate, then folds the finished exposure figure in as a disclosed dollar offset. **Part IV** sets out the two-layer attestation that lets a credentialed actuary stand behind both the methodology and any specific employer's signed figure.

HOW TO READ THIS COMPLETE BRIEF — FOUR PARTS, ONE INSTRUMENT

Part I — The Workforce Exposure Forecast (Methodology v1.0): measure marker movement, associate it with published cost literature, and project it forward with credibility weighting and disclosed confidence intervals. **Part II — Multi-Marker Cost Monetization:** the conservative, lower-bound coefficient register that converts measured marker deltas into a forward exposure figure. **Part III — The Workforce Claims Forecast / Experience Calibration (Methodology v2.1):** credibility-blend the group's own paid medical and Rx experience against a manual rate, then fold the exposure figure in as a disclosed dollar offset, returning a projected forward claims figure with two labeled uncertainty bands. **Part IV — The Two-Layer Attestation:** a methodology-level attestation statement that stands open for review, and an engagement-level Statement of Actuarial Opinion carried on each signed forecast.

DISCLOSED, NOT PROPRIETARY · ADDITIVE, NEVER A REPLACEMENT

Metra publishes its full coefficient set, credibility constants, variance model, assumption register, and projection mechanics rather than treating them as proprietary intellectual property. Reviewers are encouraged to reproduce the mechanics independently against the disclosed inputs and to report any discrepancies. The instrument is a buyer-side counterpart to the forward analytics carriers, PBMs, stop-loss writers, and reinsurers already run — it is carried into the renewal or underwriting conversation alongside those carrier-side instruments, never instead of them. Where this brief and a version-pinned methodology document differ on a number, the methodology document governs.

CONTENTS

PART I OF IV

PART I

The Workforce Exposure Forecast

Methodology v1.0. Measure workforce metabolic movement, associate it with published cost literature, and project it forward as healthcare exposure the buyer side can plan against.

This part sets out the quantitative basis for Metra's forward healthcare-exposure figures. It is deliberately explicit about what is *measured*, what is *associated* from the published literature, and what is *projected* — and about which of those steps carries uncertainty.

The central design principle is that Metra's headline exposure figure is a monetization of **movement Metra actually observes** in the enrolled workforce, not a projection from an assumed clinical efficacy. The clinical-trial literature is retained, in full, as the *association* evidence that makes measured marker improvement cost-relevant. It is not the arithmetic that produces the figure.

CORE CLAIM

1 The Actuarial Claim Metra Makes, and the Evidence It Rests On

Metra's forward exposure figure rests on a measured quantity. For every enrolled employee, the platform records actual movement in four metabolic markers over time and monetizes that movement at conservative, lower-bound cost coefficients drawn from the published literature. The figure rises and falls with what the workforce actually does. No assumed trial efficacy is applied to produce it.

The clinical-trial record — the Diabetes Prevention Program and its successors — sits one layer back. It establishes that improvement in these markers is *associated* with materially lower downstream cost at the population level. That association is what makes measured movement worth monetizing. It is evidence for the method; it is not the method itself.

> *"Metra does not project savings from an assumed intervention effect. It measures the workforce's actual metabolic movement and monetizes that movement at conservative published coefficients. The trial literature explains why the movement matters; it does not generate the figure."*

Grounded in measurement

The headline figure is a direct function of observed marker deltas, refreshed continuously as employee data changes. It is auditable against the underlying movement.

Association evidence is named

Every cost association rests on a named, peer-reviewed citation with full bibliographic form. The association is population-level and is presented as such.

Conservative coefficients

Each coefficient is the lowest defensible figure in its respective literature, not the midpoint, and the markers are stacked additively with no compounding multiplier.

Disclosed uncertainty

Metra does not claim carrier-level actuarial certainty. Forward projections are plan-level estimates subject to the credibility weighting and confidence intervals described in this part.

THE MODEL

2 Three Layers: Measure, Associate, Project

Metra's quantification separates cleanly into three layers. Keeping them distinct is what allows the headline figure to remain a measured quantity while the clinical literature does the work it is actually suited for — establishing association, not generating dollar outputs.

LAYER 01 — MEASURE

For each enrolled employee, Metra records the observed change in four metabolic markers: **weight** (lb), **systolic blood pressure** (mmHg), **fasting glucose** (mg/dL), and **waist circumference** (cm). These are observed deltas drawn from logged and device-linked inputs — continuous, employee-specific, and not assumed. They are the same four markers the METRA Score is constructed from.

LAYER 02 — ASSOCIATE

Published, peer-reviewed literature associates improvement in each of these markers with reduced downstream medical cost. This association is established at the population level — it describes expected value across a cohort, not a guaranteed outcome for any individual employee or any single plan. The clinical evidence base in the next section is this layer in full.

LAYER 03 — PROJECT

Metra projects on two distinct surfaces, and they must not be conflated:

(1) The forward exposure figure. The dashboard's Forward Healthcare Exposure figure is a direct monetization of the measured deltas: each marker delta is multiplied by its conservative coefficient and the markers are summed additively across the workforce. *No* compounding multiplier and *no* credibility adjustment is applied — it is an arithmetic monetization of what was observed.

(2) The forward forecast. Where Metra projects marker trajectories forward in time — the 30, 60, and 90-day horizons — it applies Bühlmann credibility weighting and percentile-bootstrap confidence intervals, so that smaller and sparser populations are shrunk toward the prior and every forward figure carries a disclosed interval.

THE DISTINCTION THAT MATTERS

The forward exposure figure is **not** credibility-weighted — it is a measured monetization. Credibility weighting and confidence intervals apply to the **forward forecast** of marker trajectories. A reviewer who reads the exposure figure as a credibility-adjusted projection, or the forecast as an un-intervalled point estimate, has crossed the two surfaces.

THE ASSOCIATION LAYER

3 Four Peer-Reviewed Studies Establish the Cost Association

This evidence establishes that movement in Metra's measured markers is associated with downstream cost. It is the basis for treating measured improvement as cost-relevant. It is not the projection arithmetic — the dollar figure comes from the measured-delta monetization described above, not from applying these effect sizes directly.

Study 1 — Knowler et al., NEJM 2002

58 percent reduction in T2D progression (lifestyle-intervention arm versus placebo), 3,234 participants, randomized controlled trial — the highest level of evidence. Establishes the association between structured metabolic improvement and reduced T2D progression. The effect applies to individuals with prediabetes and is a *population-level association*, retained here as evidence, not as a per-plan output.

Study 2 — Tabák et al., Lancet 2012

Measurable metabolic deterioration occurs 5 to 6 years before T2D diagnosis, with an accelerating trajectory in the final 2 to 3 years. Metra operates in the late portion of this acceleration window, where signal-to-noise on measured markers is highest.

Study 3 — Ely et al., CDC 2017

National participant-level evaluation validating digital and remote delivery of DPP lifestyle intervention as producing clinically meaningful outcomes (an associated ~40 percent risk reduction in real-world digital delivery, versus the 58 percent in-person trial figure). Provides the association basis for Metra's digitally delivered model. Outcome equivalence of Metra's AI-assisted coaching to human-coached delivery remains an assumption pending direct validation.

Study 4 — Milken Institute 2019

Cardiometabolic conditions — diabetes, cardiovascular disease, and obesity-related comorbidities — account for the largest share of US chronic-disease cost burden. Provides the population cost-attribution basis for Metra's employer financial projections. (Goetzel 2004 is retained in the citation register as supporting evidence on absence and presenteeism cost, but is no longer cited as the source for total expenditure share.)

PROJECTION — SURFACE 1

4 How Measured Movement Becomes a Dollar Figure

The forward exposure figure is produced by a single, transparent operation. For each marker, Metra multiplies the measured delta by a conservative, lower-bound cost coefficient. The markers are then summed additively across the workforce. No multiplier is applied for compounding, synergy, productivity, or absenteeism — all of which would raise the figure.

The body of this part works the example on a single marker — **weight**, at a conservative **$68 per pound per year**, the lower bound of the published Finkelstein / Tsai range. The live platform extends the identical method to the three remaining METRA-Score markers — systolic blood pressure at $150/mmHg, fasting glucose at $50/mg-dL, and waist circumference at $30/cm. The full coefficient register, with primary sources, is Part II of this brief.

MARKER	UNITS	COEFFICIENT (YR)	WORKED HERE?
Weight	per lb	**$68**	Yes — body example
Systolic blood pressure	per mmHg	**$150**	Part II register
Fasting glucose	per mg/dL	**$50**	Part II register
Waist circumference	per cm	**$30**	Part II register

The workforce total is the sum, across all participating employees, of *(measured weight loss × $68/lb) + (measured systolic drop × $150/mmHg) + (measured fasting-glucose drop × $50/mg-dL) + (measured waist reduction × $30/cm)*. Because every term is keyed to an observed delta, the figure is auditable against the underlying movement and is intended as a defensible lower-bound forward-exposure estimate for governance and renewal conversations — not a marketing figure.

PROJECTION — SURFACE 2

5 The Forward Forecast: Credibility Weighting and Confidence Intervals

Monetizing observed deltas answers "what is the exposure implied by movement to date." Projecting that movement forward answers "where is it heading." The forward forecast produces 30, 60, and 90-day projections of marker trajectories, and it is here — not in the exposure figure — that Metra applies statistical credibility weighting and discloses uncertainty as intervals. This mirrors the public Metra Methodology v1.0.

Bühlmann credibility weighting

Each plan's own observed signal is blended with a population prior using a Bühlmann credibility factor $Z = N / (N + 400)$, where N is the plan's enrolled cohort size. A large cohort earns a Z near 1 (the plan's own experience dominates); a small cohort earns a low Z and is shrunk toward the prior. This is the formal mechanism behind the small-plan limitation: below roughly 100 enrolled, Z is low and the forward intervals widen accordingly.

Participation adjustment

The enrolled cohort is participation-adjusted to an effective size $N_{eff} = N \cdot (1 - \alpha)$, with $\alpha = 0.15$, before the horizon projection is formed. The participation-adjusted horizon estimate is $E(t) = N_{eff} \cdot (Nc / N) \cdot Z \cdot \hat{\mu} \cdot (t / 12)$, where Nc/N is the participating share, $\hat{\mu}$ is the estimated annualized per-capita signal, and t is the horizon in months (so t/12 expresses the fraction of a year).

Percentile-bootstrap confidence intervals

Every forward figure carries a confidence interval generated by a percentile bootstrap of $B = 10{,}000$ resamples under a reproducible seed, so the interval is deterministic and re-derivable from the disclosed inputs. Forward projections are reported as ranges, never as bare point estimates.

AGAIN, THE TWO SURFACES ARE SEPARATE

The credibility factor Z and the bootstrap interval govern the **forward forecast** of marker trajectories. They are **not** applied to the measured exposure figure, which remains a straight additive monetization of observed deltas.

POPULATION STRATIFICATION

6 Three Actuarially Distinct Risk Cohorts

Metra's scoring engine classifies enrolled employees into three prospective risk tiers using a composite metabolic index. The distribution below reflects expected ranges for a typical US commercial employer population, anchored to NHANES 2017–2020 metabolic-syndrome prevalence. Plan-specific distribution varies by demographic profile.

Tier 1 — Improving (≈30 to 40% of population)

Low near-term claim probability, trending toward stable metabolic health. Intervention: maintenance coaching. Actuarial implication: stable or declining claims contribution over the projection period.

Tier 2 — Neutral / Elevated (≈40 to 50% of population)

Moderate metabolic risk; stable markers at an elevated baseline. Intervention: active monitoring and behavioral coaching. Actuarial implication: moderate claims growth, partially reducible with intervention.

Tier 3 — Declining (≈10 to 20% of population)

High near-term claim probability for metabolically driven conditions; accelerating deterioration consistent with the Tabák 2012 pre-diagnosis trajectory. Intervention: structured clinical program, DPP-aligned. Actuarial implication: without intervention, the primary driver of metabolically attributable trend.

Tier 3 metabolic risk overlaps with, but is not identical to, the top 10 percent of plan cost. Metra's measured exposure applies to the metabolically driven markers it observes, not to high-cost claims attributable to non-metabolic causes such as oncology, trauma, or maternity. AHRQ MEPS cost-concentration data is referenced for context, not as the basis for the exposure figure.

REFERENCE CASE

7 500-Employee Reference Case: Measured-Delta Mechanics

The following reference case illustrates the mechanics of the measured-delta monetization on the weight marker. It is an *illustration of the arithmetic*, using conservative example deltas; it is not a guarantee of outcomes for any plan. Actual figures are driven entirely by the movement Metra observes in a specific workforce.

INPUT	VALUE	BASIS
Enrolled employees	**500**	Reference case
Actively engaged cohort	**50 to 125**	10 to 25% net effective participation
Measured avg. weight loss (engaged)	**4 to 8 lb**	Conservative example delta over period
Weight coefficient	**$68 / lb / yr**	Finkelstein / Tsai lower bound
Forward exposure, weight marker only	**≈ $14K to $68K / yr**	engaged × measured lb × $68

The figure above reflects a **single marker**. The live platform stacks all four markers additively (see Part II), which raises the workforce total, and the forward forecast then projects the trajectory across the 30/60/90-day horizons with the credibility weighting and confidence intervals described earlier. Every term remains keyed to an observed delta; nothing in this case is produced by applying a clinical effect size to plan cost.

ILLUSTRATIVE CLINICAL SCENARIO

8 Association Evidence as a Scenario — Not the Projection Model

READ THIS SECTION AS ILLUSTRATION ONLY

What follows applies published DPP effect sizes to a hypothetical cohort to show *why* measured metabolic improvement plausibly reduces downstream T2D cost. It is association evidence rendered as a worked scenario. **It is not how Metra computes the dashboard's forward exposure figure** — that figure is the additive monetization of measured deltas described in the preceding sections. No plan should read the trend-point or dollar figures in this section as a Metra output.

The scenario builds a trend-modification estimate in four steps, each applying a published effect size to the 500-employee reference cohort.

Step 01 — Baseline trend

Start with the PwC HRI 2025 national employer trend (9.4 percent). Plan-specific loss ratio and prior-year claims experience replace the benchmark where available; otherwise the benchmark applies with a wider disclosed sensitivity range.

Step 02 — High-risk cohort identification

Apply METRA Score stratification to identify the Tier 3 cohort (≈10 to 20 percent of enrolled) and model its forward contribution to T2D-attributable trend.

Step 03 — Intervention effect application

Apply the DPP effect size (58 percent T2D progression reduction, Knowler 2002) to the T2D-attributable share of Tier 3 cost (estimated 25 to 40 percent, ADA 2024), adjusted by a net effective participation factor. Only avoided T2D-related claims are credited; other metabolic conditions are not.

Step 04 — Net trend modification (illustrative output)

The scenario yields a central illustrative trend modification of 1 to 2.5 percentage points below baseline, sensitivity-tested at ±0.8 points across the three highest-impact variables.

SCENARIO LINE	VALUE (ILLUSTRATIVE)	SOURCE
Base annual healthcare cost	**$9.1M**	KFF EHBS 2024, $18,200 PEPY ×500
Tier 3 cost share	**$4.6M to $5.9M**	AHRQ MEPS 2024 (context)
T2D-attributable share of Tier 3	**$1.2M to $2.4M**	ADA 2024
Central illustrative trend modification	**1 to 2.5 pts**	Applied to baseline trend
Illustrative Year 1 range	**$90K to $230K**	Scenario only — not a Metra output
Illustrative 3-year cumulative	**$370K to $940K**	Scenario only — not a Metra output

Sensitivity of the illustrative scenario

Single-variable sensitivity on the three inputs with the greatest impact on the illustrative trend figure, all using the 500-EE reference cohort.

VARIABLE	LOW	CENTRAL	HIGH
Net effective participation	8% → 0.5 pts	10–25% → 1–2.5 pts	40% → 3.4 pts
Tier 3 cohort size	8% → 0.4 pts	10–20% → 1–2.5 pts	25% → 3.1 pts
T2D share of Tier 3 cost	20% → 0.7 pts	25–40% → 1–2.5 pts	50% → 3.0 pts

Trial-level participation is unrealistic in commercial populations; the low end reflects realistic voluntary uptake at small employers without active enrollment campaigns. These figures characterize the *scenario's* sensitivity, not the measured exposure figure.

METRA SCORE METHODOLOGY

9 What the METRA Score Measures, How It Is Constructed

The METRA Score is a composite metabolic-health index, built from four input categories and updated continuously as employee data changes. The first category supplies the markers monetized in the exposure figure.

Input 01 — Primary metabolic markers (weight: high)

Fasting glucose, HbA1c (self-reported or device-linked), waist circumference, BMI, blood pressure. Source: NCEP ATP III metabolic-syndrome criteria; ADA diagnostic thresholds.

Input 02 — Behavioral risk inputs (weight: moderate)

Physical-activity frequency and intensity, dietary-pattern scoring, sleep quality, stress index. Source: validated behavioral-assessment instruments.

Input 03 — Trajectory signals (weight: high)

30, 60, and 90-day change in primary markers; directional velocity of metabolic indicators. Source: Tabák 2012 longitudinal trajectory model. These trajectory signals feed the forward forecast.

Input 04 — Engagement data (weight: low to moderate)

Coaching-session completion, milestone achievement, platform-interaction frequency. Source: DPP engagement metrics (Ely 2017) for digitally delivered intervention. Metra's AI-assisted coaching is the delivery mechanism; outcome equivalence to human-coached DPP is an assumption pending direct validation.

ASSUMPTIONS REGISTER

10 All Material Assumptions Disclosed and Auditable

The register is grouped to keep the two surfaces separate: the inputs that drive the actual measured-delta model and forward forecast, and the inputs that drive the illustrative clinical scenario only.

Group A — Measurement, monetization, and forecast inputs (the model)

INPUT	VALUE	SOURCE / LIMITATION
Weight coefficient	**$68 / lb / yr**	Finkelstein / Tsai lower bound
Systolic BP coefficient	**$150 / mmHg / yr**	Moran 2015; CDC Million Hearts (Part II)
Fasting glucose coefficient	**$50 / mg-dL / yr**	Knowler 2002; ADA (Part II)
Waist coefficient	**$30 / cm / yr**	AHRQ HCUP; Boudreau 2009 (Part II)
Stacking rule	**Additive, no multiplier**	Conservative; excludes compounding/productivity
Credibility factor (forecast)	**Z = N / (N + 400)**	Bühlmann; small-N shrinks to prior
Participation adjustment (forecast)	**N_eff = N · (1 − 0.15)**	Applied before horizon projection
Confidence interval (forecast)	**Percentile bootstrap, B = 10,000**	Reproducible seed; ranges not point estimates
Forecast horizons	**30 / 60 / 90 days**	Tabák 2012 trajectory window

Group B — Illustrative clinical scenario inputs (not the model)

INPUT	VALUE	SOURCE / LIMITATION
Baseline trend	**9.4%**	PwC HRI 2025; national benchmark
Tier 3 cohort size	**10 to 20% enrolled**	NHANES 2017–2020; demographic variation
T2D-attributable share of Tier 3	**25 to 40%**	ADA 2024; plan disease mix shifts share
T2D prevention effect size	**58% reduction**	Knowler 2002; RCT, prediabetic individuals
Net effective participation	**10 to 25% of Tier 3**	enrollment × completion; modeled separately
Detection window	**12 to 24 months pre-diagnosis**	Tabák 2012 late-acceleration phase

INPUT	VALUE	SOURCE / LIMITATION
T2D annual cost premium	**$9,601 to $12,022**	ADA Diabetes Care 2024

DATA GOVERNANCE & PRIVACY

11 De-Identified Inputs and Aggregate Reporting Controls

Metra is designed around privacy-by-design principles and works with de-identified, aggregate, or otherwise non-PHI data for actuarial and compliance review.

Data architecture

Metra ingests de-identified population-level data and non-PHI operational inputs, organized to support cohort analysis, trend monitoring, and actuarial review. Access is structured by role and purpose to limit exposure to sensitive information.

De-identification standard

Employer-facing reporting uses aggregated outputs with minimum cohort thresholds. Small or sparse groups are suppressed to reduce re-identification risk. Individual-level outputs are not exposed in employer reporting views.

Actuarial data access

Actuaries may review aggregate population data through governed data-sharing workflows. Carrier and employer analyses rely on summary-level outputs, not individual records. Any research or validation linkage requires a separate approved process and scoped data access.

VALIDATION PLAN

12 How Metra Validates Projected Versus Realized Outcomes

Metra treats the forward projection as a falsifiable prediction, not a marketing claim. The validation pathway has three components, each comparing projected against realized outcomes on a defined cadence.

Component 01 — Continuous cohort outcome tracking

For every deployment, Metra tracks projected versus realized cohort outcomes at 6, 12, and 24-month intervals — Tier 3 progression rate, intervention completion, biometric trajectory by tier, and metabolically attributable claims emergence where claims data is shared. Discrepancies inform iterative model calibration quarterly.

Component 02 — Annual retrospective trend comparison

Once per plan year, Metra publishes a retrospective comparison of projected versus observed plan trend, controlling for plan-design changes (deductible, formulary, network) that would otherwise contaminate the comparison. Shared with the contracting employer and available in summary form for actuarial review under governed access.

Component 03 — Independent actuarial replication

An independent actuarial replication study is being scoped for the second full deployment year: an outside actuary, contracted at arm's length, reproduces the projection mechanics on a deployed cohort using the disclosed register and reports any discrepancies. Findings will be published in summary form regardless of outcome.

> *"A model unwilling to be tested against its own predictions is not a model. It is a marketing claim. The validation pathway is the structural commitment to the former and against the latter."*

MODEL LIMITATIONS

13 What Metra Does Not Claim — Exposure Forecast Boundaries

The following limitations are disclosed in full and should be taken into account when evaluating Metra's projections for plan pricing, reserve setting, or renewal-negotiation support.

Individual clinical outcomes

Metra makes no clinical predictions about individual employees. The METRA Score is a risk-stratification index, not a diagnostic instrument, and is not a substitute for clinical evaluation.

Carrier-level actuarial certification

Metra's projections are plan-level estimates, not actuarially certified trend indications, and are not a substitute for carrier underwriting or a regulatory actuarial opinion.

Association, not causal guarantee

The cost associations underlying the exposure figure reflect expected value across a population, not a guaranteed outcome for any specific plan. They are applications of published, population-level associations, not causal attributions.

Participation assumption

The illustrative scenario assumes net effective participation of 10 to 25 percent among the Tier 3 cohort (enrollment 15–30% × completion 70–80%). The measured exposure figure carries no such assumption — it monetizes only the movement actually observed — but lower engagement means fewer measured deltas to monetize.

Small-plan credibility

Plans below roughly 100 enrolled employees receive a low Bühlmann credibility factor ($Z = N/(N+400)$) on the forward forecast and therefore wider confidence intervals. The stratification model requires a minimum enrolled population to generate statistically stable cohort estimates.

Data completeness

Exposure figures and METRA Scores are more accurate with more complete self-reported and device-linked inputs. Employees who do not fully complete onboarding contribute lower-confidence data.

Non-metabolic cost categories

The exposure figure is restricted to the metabolic markers Metra measures. Claims attributable to oncology, trauma, maternity, mental health, and other non-metabolic categories are not credited, even within the same high-risk cohort. Plans where non-metabolic high-cost claims dominate will see proportionally smaller measured exposure.

PART II

Multi-Marker Cost Monetization

The conservative, lower-bound coefficient register that converts each measured marker delta into dollars — extending the body's single-marker weight example to the full METRA-Score marker set.

COEFFICIENT REGISTER

14 The Multi-Marker Cost Coefficient Register

Part I monetizes a single biometric marker — weight — at a conservative $68/lb/yr. That figure remains the lower bound of the published Finkelstein / Tsai range and is unchanged. This register extends the same conservative-anchor methodology to the three remaining METRA-Score components — systolic blood pressure, fasting glucose, and waist circumference — so the dashboard's forward exposure figure reflects the workforce's actual movement across every marker the score is built from, not weight alone.

MARKER	UNITS	COEFFICIENT (YR)	PRIMARY SOURCE
Weight	per lb	**$68**	Finkelstein et al., *Annual Medical Spending Attributable to Obesity* (Health Affairs, 2009; 2014 update); Tsai et al., *Obesity Reviews* (2011). Lower bound of insured working-age range.
Blood pressure	per mmHg systolic	**$150**	Moran et al., *Annals of Internal Medicine* (2015) — Cost-Effectiveness of Hypertension Therapy; CDC Million Hearts Initiative. Lower bound of CV-event-avoided cost range.
Fasting glucose	per mg/dL	**$50**	Knowler et al., *NEJM* (2002) — Diabetes Prevention Program; ADA, *Economic Costs of Diabetes in the U.S.* (2017 update). Derived from per-mg/dL share of total cost of care avoided.
Waist	per cm	**$30**	AHRQ HCUP — Statistical Briefs on metabolic-syndrome inpatient/outpatient cost; Boudreau et al., *Population Health Management* (2009). Lower bound; waist is the strongest single anthropometric predictor of metabolic-syndrome-driven medical spend.

CONSERVATISM & STACKING

Each coefficient is the lowest defensible figure in its respective literature, not the midpoint. The four markers are stacked additively against each employee's actual measured deltas: the workforce total is the sum of *(avg weight loss × $68/lb) + (avg systolic drop × $150/mmHg) + (avg fasting glucose drop × $50/mg/dL) + (avg waist reduction × $30/cm)*. No multiplier is applied for compounding, synergistic effects, downstream productivity, or absenteeism — all of which would raise the figure. The result is intended as a defensible lower-bound forward-exposure estimate suitable for governance and renewal conversations, not a marketing figure.

PART III

The Workforce Claims Forecast

Methodology v2.1 — Experience Calibration. The renewal math a carrier already runs on the employer's population, run by the buyer side, with Metra's forward metabolic signal folded in as a disclosed, auditable dollar offset.

This part sets out the quantitative basis for Metra's *projected forward claims* figure. It builds directly on Parts I and II: it takes the finished forward exposure figure and does with it what a carrier already does at renewal — nets it against a credibility-blended projection of the group's own paid claims.

The central design principle carries forward 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 V2.1 CHANGED FROM V2.0

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 PART — 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 Part I — returning a projected claims figure with two labeled uncertainty bands.

CORE CLAIM

15 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

16 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 Part I 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

17 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

18 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 $k = 400$ 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 $Z^* = 0.90$, reached at $N^* = 3{,}600$ 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 Part I'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 k = 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, $a = (\tau \cdot E)^2$, the constant implies a per-life process SD of $s = \sqrt{(k \cdot a)} = \sqrt{k} \cdot \tau \cdot E$ — an implied per-life pooled-claims CV of $\sqrt{k} \cdot \tau = \sqrt{400} \cdot 0.15 =$ **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

19 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 Part I 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 Part I 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 Part I 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

20 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	Part III, Layer 02
$Z_{metabolic}$	The cohort biometric exposure signal	A conservative prior	Part I, inside the offset M

The metabolic offset M arrives already credibility-weighted by Part I'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

21 322-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 Part I 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

22 Two Labeled Bands, One Variance Model

v2.0 inherited its band entirely from the Part I 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 *estimated mean* and the volatility of *realized* 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 $a = (\tau \cdot E)^2$ and the per-life process SD is $s = \sqrt{(k \cdot a)} = \sqrt{k} \cdot \tau \cdot E$.

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

This is the estimation uncertainty of the *mean* 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 *realized* 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 $\hat{G}$ 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

23 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 Part I 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

24 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

25 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 Part I and may double-count across markers

Part I 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 *reduces* 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 Part I'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.

PART IV

The Two-Layer Attestation

How a credentialed actuary stands behind the instrument: a methodology-level statement that stands open for review, and an engagement-level Statement of Actuarial Opinion carried on each signed forecast.

The two analytic surfaces above are only as useful as the attestation discipline around them. Metra separates that discipline into two layers, deliberately, so that a reviewing actuary can endorse the *methodology* without being made to appear to endorse any particular employer's number, and can opine on a particular employer's number without being made to re-derive the methodology each time.

LAYER 1 — METHODOLOGY

26 Layer 1 — The Methodology-Level Attestation

Metra publishes its full methodology corpus — the version-pinned specifications behind Parts I through III — precisely so a credentialed actuary can review the methods and, if they concur, attest at the methodology level. This statement is framed as available for attestation and standing open for review; it is not a claim that any particular actuary has already signed it.

METHODOLOGY-LEVEL ATTESTATION STATEMENT

A credentialed actuary reviewing Metra Methodology v2.1 may attest that, in their opinion, the credibility procedure (limited-fluctuation full-credibility anchor, $Z^* = 0.90$ at $N^* = 3{,}600$ life-years ⇒ $k = 400$, applied via Bühlmann–Straub), the two-band variance model (an expected-claims confidence band and a realized-paid predictive interval from a single disclosed dispersion input τ), and the disclosed coefficients and limitations are reasonable and appropriate for producing a buyer-side forward estimate of workforce claims exposure, and that the methods and communications conform to ASOP Nos. 25, 41, and 56. This methodology-level statement does not opine on any specific employer's figure — that is the engagement-level Statement of Actuarial Opinion carried on each signed Workforce Claims Forecast, applied to a named group's connected experience (ASOP No. 23) and version-pinned to v2.1. Together they form the buyer-side instrument's two-layer attestation: the methodology stands open for review, and each signed forecast carries a defined-scope opinion on the group's own number.

What the methodology-level statement covers

The reasonableness and appropriateness of the credibility procedure, the variance model, the disclosed coefficients, and the stated limitations — and the conformance of the methods and communications with ASOP Nos. 25, 41, and 56. It is an opinion about the *method*, applicable across every group the method is run on.

What it deliberately does not cover

It does not opine on the data quality, assumptions, or resulting figure for any specific employer. A group's completion factor, pooling point, manual rate, and connected experience are engagement facts — they belong to Layer 2 under ASOP No. 23, not to a statement about the method.

LAYER 2 — ENGAGEMENT

27 Layer 2 — The Engagement-Level Statement of Actuarial Opinion

Each signed Workforce Claims Forecast carries a defined-scope Statement of Actuarial Opinion on its own figures. Where Layer 1 endorses the method in general, Layer 2 is the named, credentialed actuary's opinion on *this* group's number — scoped to the document in hand, version-pinned to the methodology version that produced it.

Rendered only for a real, version-matched, signed tier

The opinion is rendered live only when the forecast carries a genuine, version-matched attestation from a signing — or broker-of-record — actuary at a signed tier. A demonstration document is labelled explicitly as illustrative and carries no personal opinion; an unsigned document is methodology-only and states plainly that no engagement-level opinion attaches. The signed PDF and its cover cannot diverge on this: a single posture governs both, so a document can never show "signed" on one surface and "unsigned" on another.

What the live opinion says, and what it withholds

The live opinion names the actuary, scopes itself to the figures in that specific document, identifies the methodology version, and invokes ASOP Nos. 23, 25, 41, and 56 — data quality on the group's connected experience, the credibility procedure, the communication of assumptions and both uncertainty bands, and the model. It closes as it must: as a buyer-side estimate carried alongside the carrier-side instruments, **not** a certification of a rate.

INTEGRITY DISCIPLINE

28 How the Two Layers Interlock — and the Version-Pinning Discipline

The two layers are designed to be complementary and non-substitutable. The methodology stands open for review at Layer 1; each signed forecast carries a defined-scope opinion on the group's own number at Layer 2. Neither layer can be made to do the other's work: a methodology endorsement is never presented as an opinion on a particular employer, and an engagement opinion never silently re-certifies the method beyond the version it was written against.

PROPERTY	LAYER 1 — METHODOLOGY	LAYER 2 — ENGAGEMENT
Opines on	The method, across all groups	One named group's own figure
Primary ASOPs	Nos. 25, 41, 56	Nos. 23, 25, 41, 56
Where it lives	The public methodology corpus	The signed Workforce Claims Forecast PDF
Posture	Available for attestation; stands open for review	Live only for a real, version-matched, signed tier

Why the corpus is version-pinned

Attestations do not transfer across methodology versions. Each version — v1.0, v2.0, v2.1 — stays published and version-pinned, and a signed opinion cites exactly the version that generated the figure. This is what makes both layers honest: a reader can always tell which methodology version a given signature actually certified, and superseding a version requires a fresh attestation rather than silently inheriting the old one. The cost of that discipline is that the corpus is divided by version — which is why this complete edition presents each version as a labeled part rather than rewriting them into a single moving target.

One instrument, two instruments' worth of carriers' science, carried buyer-side

Both analytic surfaces and both attestation layers serve a single purpose: to give the buyer side the forward analytic capability the carrier already has. The exposure forecast and the claims calibration are buyer-side instruments, disclosed in full and carried into the renewal or

underwriting conversation alongside the carrier-side instruments — never instead of them. Neither certifies a rate; together they let the CFO plan, the CHRO govern, the broker carry the number into the carrier conversation, and the actuary fold it into their own credibility math.

FULL CITATION REGISTER — ALL PARTS, FULL BIBLIOGRAPHIC FORM

1. Knowler WC, Barrett-Connor E, Fowler SE, et al. Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin. **N Engl J Med.** 2002;346(6):393–403.
2. Tabák AG, Herder C, Rathmann W, Brunner EJ, Kivimäki M. Prediabetes: a high-risk state for diabetes development. **Lancet.** 2012;379(9833):2279–2290.
3. Ely EK, Gruss SM, Luman ET, et al. A national effort to prevent type 2 diabetes: participant-level evaluation of CDC's National Diabetes Prevention Program. **Diabetes Care.** 2017;40(10):1331–1341.
4. Goetzel RZ, Long SR, Ozminkowski RJ, Hawkins K, Wang S, Lynch W. Health, absence, disability, and presenteeism cost estimates of certain physical and mental health conditions affecting U.S. employers. **J Occup Environ Med.** 2004;46(4):398–412.
5. American Diabetes Association. Economic costs of diabetes in the U.S. in 2022. **Diabetes Care.** 2024;47(1):26–43.
6. Finkelstein EA, Trogdon JG, Cohen JW, Dietz W. Annual medical spending attributable to obesity: payer- and service-specific estimates. **Health Aff.** 2009;28(5):w822–w831 (2014 update).
7. Tsai AG, Williamson DF, Glick HA. Direct medical cost of overweight and obesity in the USA: a quantitative systematic review. **Obes Rev.** 2011;12(1):50–61.
8. Moran AE, Odden MC, Thanataveerat A, et al. Cost-effectiveness of hypertension therapy according to 2014 guidelines. **Ann Intern Med.** 2015 (with CDC Million Hearts Initiative).
9. Boudreau DM, et al. Health care utilization and costs by metabolic-syndrome risk factors. **Popul Health Manag.** 2009;12(1):37–45.
10. Agency for Healthcare Research and Quality (AHRQ). Medical Expenditure Panel Survey (MEPS) and HCUP Statistical Briefs on metabolic-syndrome cost. 2024.
11. KFF (Kaiser Family Foundation). Employer Health Benefits Survey 2024. Section 1: Cost of Health Insurance.
12. PricewaterhouseCoopers Health Research Institute. Medical Cost Trend: Behind the Numbers 2025. PwC HRI; 2024.
13. Bureau of Labor Statistics. Employer Costs for Employee Compensation, September 2024. USDL-24-2330.
14. National Cholesterol Education Program (NCEP) Expert Panel. Third Report (ATP III). NIH Publication No. 02-5215. 2002.
15. U.S. Department of Health and Human Services. HIPAA Administrative Simplification: Regulation Text. 45 CFR Parts 160, 162, and 164. 2013 Omnibus Final Rule.
16. Centers for Disease Control and Prevention. National Diabetes Prevention Program. Program description and eligibility criteria. cdc.gov/diabetes/prevention. Accessed 2024.
17. Milken Institute. The Cost of Chronic Diseases in the U.S. Milken Institute, 2018 (updated 2019).
18. Centers for Disease Control and Prevention. National Health and Nutrition Examination Survey (NHANES), 2017–2020 cycle. Metabolic syndrome prevalence estimates.
19. Bühlmann H. Experience rating and credibility. **ASTIN Bulletin.** 1967;4(3):199–207.

20. Bühlmann H, Straub E. Glaubwürdigkeit für Schadensätze. **Bulletin of the Swiss Association of Actuaries.** 1970.
21. Longley-Cook LH. An introduction to credibility theory. **Proceedings of the Casualty Actuarial Society.** 1962 (limited-fluctuation full-credibility standards).
22. Herzog TN. **Introduction to Credibility Theory.** ACTEX Publications.
23. Klugman SA, Panjer HH, Willmot GE. **Loss Models: From Data to Decisions.** 5th ed. Wiley; 2019.
24. Actuarial Standards Board. **ASOP No. 23**, Data Quality.
25. Actuarial Standards Board. **ASOP No. 25**, Credibility Procedures.
26. Actuarial Standards Board. **ASOP No. 41**, Actuarial Communications.
27. Actuarial Standards Board. **ASOP No. 56**, Modeling.
28. Society of Actuaries. Group health renewal and manual-rating practice notes.
29. Metra Healthcare Intelligence. **Methodology v1.0 — Workforce Exposure Forecast** and **Methodology v2.1 — Experience Calibration** (additive to v2.0). usemetra.com/methodology/.