

METHODOLOGY · V1.1 · WORKFORCE EXPOSURE FORECAST · FOR REVIEW & ATTESTATION

Metra Workforce Exposure Forecast Methodology v1.1

Public technical specification for the Metra Workforce Exposure Forecast engine. v1.1 replaces the four-marker additive dollar stacking of v1.0 with a single BMI economic pathway, a versioned coefficient register, explicit eligibility rules, and a deterministic digest-seeded bootstrap.

Version 1.1-draft **Register** 2026-08-refresh

Coefficient status PROPOSED — Pending actuarial review

Layer Exposure (Workforce Exposure Forecast)

Companion Methodology v2.1 — Experience Calibration

CONTENTS

1. Scope and intent
2. What changed from v1.0
3. Notation and definitions
4. Population and denominators
5. Versioned coefficient register
6. BMI economic pathway
7. Point estimate
8. Deterministic bootstrap confidence interval
9. Bidirectional marker visibility
10. HbA1c conditional module (disabled)
11. Interpretation and boundary language

- 12. Limitations
- 13. Attestation protocol
- 14. Version history
- 15. Bibliography

§1 Scope and intent

This document specifies the Metra Workforce Exposure Forecast, Methodology v1.1. The Forecast is a **buyer-side, forward-looking, biometric-anchored economic signal** intended to be carried into renewal, underwriting, and governance conversations alongside — never instead of — carrier-side claim experience instruments.

The Forecast estimates, in annualised dollar terms, the healthcare-expenditure difference associated with the qualifying *measured* BMI movement in a given employer cohort. It is not a claim prediction, not a premium projection, and not a guaranteed savings figure. See §11 and §12 for the full boundary language.

What this methodology specifies:

- The versioned coefficient register (§5) — every economic parameter the engine may use, with status, eligibility constraints, and evidence source
- The BMI economic pathway (§6) — the single active economic module for v1.1
- How the point estimate is computed from qualifying participants (§7)
- The deterministic bootstrap confidence interval (§8)
- The bidirectional marker visibility rule (§9)
- The HbA1c conditional module, currently disabled (§10)
- The interpretation boundary language required on every output (§11)
- Limitations named by Metra (§12)
- The attestation protocol (§13)

What this methodology does not specify: the v2.1 Experience Calibration addendum (claims-layer credibility blend, forward claims projection, two-band confidence/predictive interval). That is specified separately in [Methodology v2.1](#), which consumes the v1.1 Exposure point estimate as a finished input.

§2 What changed from v1.0

v1.0 monetised four markers independently and additively: weight (\$68/lb), systolic blood pressure (\$150/mmHg), fasting glucose (\$50/mg-dL), and waist circumference (\$30/cm). The coefficients were applied to every participant's measured delta, summed, and the total reported as the workforce's forward exposure estimate. A Bühlmann credibility blend ($Z = N/(N+400)$), a lag haircut ($\lambda = 0.50$), and an attrition-horizon projection were applied to the dollar total.

v1.1 replaces this entirely. The 2026 Evidence Refresh found that defensible linear \$/unit associations no longer exist for the per-mmHg, per-mg/dL, and per-cm markers in the contemporary peer-reviewed literature — the associations are categorical (hypertension cost bands, diabetic subgroups) rather than continuous \$/unit, making the v1.0 linear coefficients unjustifiable. The weight-per-lb coefficient was also retired: BMI is a more defensible population-level anchor than per-pound weight, and the Thorpe & Joski (2024) publication provides a direct \$/BMI-point association with a citable lower 95% CI bound.

Specific changes in v1.1:

- **BMI replaces weight-per-lb** — the active economic pathway is \$93/BMI-point/person-year, the lower 95% CI bound of the Thorpe & Joski (2024) \$/BMI-point-above-30 spending association for ESI adults. Status: PROPOSED pending actuarial review.
- **SBP, fasting glucose, waist — non-monetised** — these markers remain in the instrument as risk-intelligence signals (trajectory, movement direction) but contribute zero dollars under v1.1. The legacy \$150/mmHg, \$50/mg-dL, and \$30/cm coefficients are retired.
- **No Bühlmann credibility blend in the exposure engine** — v1.1 has no credibility weight applied to the exposure point estimate. Credibility blending exists in the separate v2.1 claims calibration layer. (The `credibility_weight` database column retains a stored value of 1.0 for v1.1 rows because the column has a NOT NULL constraint on the underlying table; this is not a calculation.)
- **No lag haircut** — v1.1 does not apply a lag adjustment to the dollar output. The per-person-year coefficient is applied to measured BMI movement directly.

- **No horizon/attrition projection** — the point estimate is an annual figure by the coefficient's units (USD per BMI-point per person-year). No attrition or horizon projection is applied within the exposure engine. Attrition scaling exists only inside the v2.1 claims calibration step.
- **Deterministic bootstrap** — the confidence interval seed is now derived deterministically from a SHA-256 digest of the canonical participant snapshot, not from a random seed. This makes the CI reproducible bit-for-bit from the same data.
- **Versioned coefficient register** — every coefficient the engine may use is recorded in a JSON register with its status (active / proposed / disabled / non_monetised), eligibility rules, evidence ID, and effective date. No coefficient may contribute dollars unless it is `active` in the register and passes eligibility.
- **Explicit eligibility rules** — the BMI pathway has codified eligibility (age 24–64, ESI enrolled, baseline BMI 30.0–45.0, current BMI ≥ 30.0 , measurement window ≥ 30 days). Participants outside these bounds receive zero credit and are excluded from the denominator for the economic estimate.
- **HbA1c conditional module** — a separate HbA1c pathway is specified in the register at \$545 per 1.0 percentage-point reduction (T2D subgroup, Lage & Boye 2020) but is `disabled` by default. It cannot be enabled without an explicit methodology version bump and independent actuarial acceptance.
- **Bidirectional visibility** — the risk-intelligence layer surfaces worsening-marker movement explicitly (it is never silent). Worsening contributes zero dollars to the economic estimate but is always reported.

§3 Notation and definitions

Symbol / term	Definition
i	Index over participants in the measured population
$\text{BMI}_{\text{baseline},i}$	Participant i 's BMI at first weight log within the measurement window (kg/m^2), rounded to 2 decimal places

Symbol / term	Definition
$BMI_{current,i}$	Participant i 's BMI at most recent weight log (kg/m ²), rounded to 2 decimal places. Must be ≥ 30 days after baseline log.
d_i	Credited BMI reduction = $\max(0, BMI_{baseline,i} - \max(BMI_{current,i}, 30))$
c	Active BMI coefficient from the register = 93.0 USD/BMI-point/person-year
S_i	Annual associated exposure shift for participant $i = d_i \times c$
$\hat{E}$	Workforce point estimate = $\sum S_i$ over qualified participants
B	Number of bootstrap resamples = 10,000
CI_{95}	Percentile bootstrap 95% confidence interval $[l_o, h_l]$
ESI	Employer-sponsored insurance. Proxied by enrollment in the Metra platform (ASSUME_ESI_FOR_ENROLLED).
$N_{enrolled}$	Total enrolled employees in scope
$N_{measured}$	Participants with valid first + last weight logs ≥ 30 days apart, plausible anthropometric values
$N_{qualified}$	Measured participants who pass all BMI-pathway eligibility rules

§4 Population and denominators

Every Workforce Exposure Forecast output reports an explicit denominator table. Each quantity is distinct and must not be conflated:

Quantity	What it counts	Used for
$N_{enrolled}$	Total enrolled employees in the company (or scoped org unit)	Participation rate denominator; representation context

Quantity	What it counts	Used for
N_{measured}	Employees with plausible, paired weight logs ≥ 30 days apart, valid age and height	Data-coverage context
$N_{\text{qualified}}$	Measured employees who pass all BMI-pathway eligibility rules (see §6.1)	Economic estimate denominator; bootstrap input
Data gaps	Missing height/age, short measurement windows, implausible anthropometrics — flagged, never imputed	Transparency; data-quality review

No extrapolation to nonparticipants. The economic estimate is computed only from qualifying measured participants. The forecast makes no projection from participants to the broader enrolled or eligible population.

ESI proxy. The platform currently does not store an employer-sponsored-insurance flag at the employee level. ESI eligibility is proxied by enrollment in the Metra platform (all enrolled employees are treated as ESI). This is a named limitation (L1).

Data validation. Participants are excluded from the measured population — never silently imputed — if they are missing height or age, have only a single weight log, have a measurement window shorter than 30 days (the register-defined minimum), or have implausible anthropometric values (height outside 120–220 cm, weight outside 30–300 kg). Exclusion counts are reported explicitly.

§5 Versioned coefficient register

Every coefficient the v1.1 engine may use is recorded in a JSON register (`coefficient-register.v1_1.json`, register version `2026-08-refresh`). No coefficient may contribute dollars to the forecast unless its status in the register is `active`. The register enforces that exactly one economic pathway is active at any time.

Status lifecycle:

- **ACTIVE** — parameter is in production use and contributes to the economic estimate
- **PROPOSED — PENDING ACTUARIAL REVIEW** — parameter is implemented and used in production but has not yet received independent actuarial sign-off; every output carrying a PROPOSED parameter discloses this status
- **DISABLED — PENDING ACCEPTANCE** — module is implemented in the register but suppressed; enabling requires an explicit methodology version bump and independent actuarial review
- **NON_MONETISED** — marker is tracked for risk intelligence only; contributes zero dollars; legacy linear coefficients for this marker are retired
- **SUPERSEDED** — retired coefficient from a prior methodology version; no longer active

ID	Marker	Value	Units	Status
BMI_ESI_2426_64_THORPE2024_LB93	BMI	\$93.00	USD/BMI-point/person-year	PROPOSED — PENDING ACTUARIAL REVIEW
HBA1C_T2D_LAGE2020_545	HbA1c	\$545.00	USD/1.0 ppt reduction/person-year	DISABLED — PENDING ACCEPTANCE
SBP_NON_MONETISED	Systolic BP	—	Risk intelligence only	NON_MONETISED

ID	Marker	Value	Units	Status
FASTING_GLUCOSE_NON_MONETISED	Fasting glucose	—	Risk intelligence only	NON_M
WAIST_NON_MONETISED	Waist circumference	—	Risk intelligence only	NON_M

ID	Marker	Value	Units	Status
weight-per-lb (\$68/lb)	Weight	—	—	SUPERSEDED (v1.0)

The register is served as an audit view at `GET /api/corporate/exposure-audit` (company admin / reviewer access). The audit view exposes coefficient IDs, statuses, eligibility rules, evidence IDs, population denominators, data-gap counts, snapshot digest, bootstrap seed, and the boundary interpretation string. It contains no employee identity.

§6 BMI economic pathway

§6.1 Eligibility rules

A participant is *economically qualified* for the BMI pathway if and only if all of the following hold (all rules are register-driven and may change only with a version bump):

- **ESI enrolled** — enrolled in an employer-sponsored health plan (proxied by Metra enrollment; see L1)
- **Age 24-64** — age at measurement ≥ 24 and ≤ 64 (inclusive)
- **Baseline BMI 30.0-45.0** — BMI at the first weight log ≥ 30.0 and ≤ 45.0 (inclusive). Participants below 30 were not obese at baseline; participants above 45 are outside the evidence boundary. No extrapolation to either group.
- **Current BMI ≥ 30.0** — BMI at the most recent weight log ≥ 30.0 . The credit floor is BMI = 30; improvement below 30 does not generate additional credit.
- **Measurement window ≥ 30 days** — the time between the first and most recent weight log must be at least 30 calendar days (register parameter `min_measurement_window_days`). Participants with shorter windows are excluded and counted in data gaps.

Participants who do not meet all five rules receive zero credit and are excluded from $N_{\text{qualified}}$. They remain in N_{measured} for transparency.

§6.2 Credit formula

CREDITED BMI REDUCTION PER PARTICIPANT

$$d_i = \max(0, \text{baseline_bmi}_i - \max(\text{current_bmi}_i, 30))$$

The inner $\max(\cdot, 30)$ floors the effective current BMI at 30.0, so movement below the obesity threshold generates no additional credit. The outer $\max(0, \cdot)$ ensures that BMI deterioration (increase) contributes zero dollars — it is surfaced by the risk-intelligence layer (§9) but never charged to the economic estimate.

§6.3 Worked example

Suppose participant A: age 52, ESI-enrolled, baseline BMI 36.4, current BMI 33.1, measurement window 180 days.

- Eligibility: age ✓, ESI ✓, baseline 30–45 ✓, current ≥ 30 ✓, window ≥ 30 d ✓ → qualified
- $d = \max(0, 36.4 - \max(33.1, 30)) = \max(0, 36.4 - 33.1) = \mathbf{3.3 \text{ BMI points}}$
- $S = 3.3 \times \$93 = \mathbf{\$306.90/\text{person-year}}$

Suppose participant B: age 52, ESI-enrolled, baseline BMI 36.4, current BMI 38.2 (deterioration).

- Eligibility: passes ✓ (current BMI ≥ 30 , baseline in range)
- $d = \max(0, 36.4 - \max(38.2, 30)) = \max(0, 36.4 - 38.2) = \max(0, -1.8) = \mathbf{0}$
- $S = \$0.00$ — deterioration surfaced as worsening by risk-intelligence layer, not charged

Suppose participant C: age 52, ESI-enrolled, baseline BMI 36.4, current BMI 28.5 (improved past floor).

- Eligibility: current BMI $28.5 < 30 \rightarrow \mathbf{\text{not qualified}}$; excluded from $N_{\text{qualified}}$

- Credit = 0 (outside eligibility)

§6.4 Evidence basis for the \$93 coefficient

The coefficient \$93/BMI-point/person-year is the lower bound of the 95% confidence interval around the association between BMI above 30 and annual employer-sponsored healthcare spending, reported by Thorpe & Joski (2024, *JAMA Network Open*). The central estimate in Thorpe & Joski (2024) is approximately \$326/BMI-point. Metra adopts the lower 95% CI bound as the most conservative defensible anchor.

Important: the Thorpe & Joski association is observational, not causal. The coefficient represents the per-BMI-point annual spending difference between obese and healthy-weight ESI employees — it does not assert that one BMI-point of measured improvement causes exactly \$93 in claim reduction. The boundary interpretation required on every output reflects this (§11).

The coefficient status is **PROPOSED — PENDING ACTUARIAL REVIEW**. Independent review by a credentialed actuary (ASA, FSA, MAAA, FCAS, or ACAS) has not yet concluded. Until that review is complete and recorded on the attestation roster, every Forecast generated from v1.1 carries the PROPOSED disclosure.

§7 Point estimate

WORKFORCE POINT ESTIMATE

$$\hat{E} = \sum d_i \times c \quad (\text{sum over economically qualified participants } i)$$

where $c = 93.0$ USD/BMI-point/person-year from the active register entry.

The point estimate is a simple sum — no credibility blend, no lag haircut, no attrition adjustment, no horizon projection. Each qualified participant's credited BMI reduction is multiplied by the register coefficient and the products are summed. The coefficient's units (per person-year) make the result an annualised figure directly.

When $N_{\text{qualified}} = 0$, the point estimate is \$0 and no bootstrap CI is computed. The Forecast reports this explicitly rather than fabricating a non-zero figure.

§8 Deterministic bootstrap confidence interval

§8.1 Procedure

The bootstrap CI characterises sampling uncertainty over the qualified population's credited BMI reductions. The procedure is deterministic: the same input population always produces the same CI.

1. **Canonical snapshot.** Serialise the participants *in caller order* — no sorting is applied. Each participant is serialised as a Python-repr JSON object: `{"age":<int>, "baseline_bmi":<float>, "current_bmi":<float>, "employer_sponsored_insurance":<0|1>, "participant_id":<str>}`. Floats use Python shortest-repr (integer-valued floats render as `30.0`, not `30`). Concatenate as a JSON array with no extra whitespace. *Because the digest is order-dependent, any implementation that reorders participants relative to the caller's input will produce a different digest, seed, and CI — use the same input order to reproduce the acceptance-gate values.*
2. **Digest.** Compute SHA-256 of the canonical JSON string (UTF-8 encoded). Record the hex digest as the snapshot fingerprint.
3. **Seed derivation.** Take the first 16 hexadecimal characters of the digest. Parse as a 64-bit unsigned integer (big-endian hex). This is the bootstrap seed.
4. **RNG.** Initialise a CPython-compatible MT19937 Mersenne Twister from the 64-bit seed using the Python `init_by_array` path (the seed is split into two 32-bit words). All subsequent draws use CPython's `getrandbits(32)`-based floating-point generation, exactly matching `random.choice()` from the reference implementation.
5. **Resampling.** Let $n = N_{\text{qualified}}$ and $Q = [d_1, \dots, d_n]$ the vector of credited reductions. For $b = 1 \dots B$ ($B = 10,000$): draw n values with replacement from Q using the RNG; accumulate the sum; multiply by the coefficient c . Store the result as T_b .
6. **Percentile CI.** Sort $T_{1 \dots B}$. The 95% interval is:

- $lo = T_{\text{sorted}}[\lfloor 0.025 \times B \rfloor]$
- $hi = T_{\text{sorted}}[\lfloor 0.975 \times B \rfloor - 1]$

7. **Null CI.** If $N_{\text{qualified}} = 0$, no CI is computed. The output field is null, never zero.

§8.2 Acceptance test

The implementation is validated against the following reference values derived from the supplied synthetic dataset:

SYNTHETIC DATASET ACCEPTANCE GATE

- Qualified participants: **45**
- Point estimate: **\$6,052.44**
- 95% CI: **[\$5,142.90, \$6,968.49]**
- Snapshot digest: **1f98b86c...** (first 8 hex chars)

The TypeScript engine matches the Python reference implementation bit-for-bit on this dataset. Any future implementation must reproduce all four values to be considered conformant.

§8.3 What the CI does and does not cover

The bootstrap CI reflects *sampling uncertainty* — the variability in the point estimate due to resampling the observed qualified cohort. It does not cover:

- Coefficient uncertainty (the \$93 figure is a point value from the register, not a distribution)
- Future participant attrition or engagement change
- Model uncertainty in the observational Thorpe & Joski association
- ESI-proxy error (enrolled $\neq$ ESI is an unquantified gap)

These are named limitations (§12).

§9 Bidirectional marker visibility (never-silent rule)

The v1.1 instrument surfaces *all* measured marker movement — improvement and deterioration — in the risk-intelligence layer. This is not optional and may not be suppressed or relabelled as neutral.

Bidirectional movement is reported for all tracked markers (BMI, systolic BP, fasting glucose, waist circumference): improving, deteriorating, and stable counts are reported separately, with the count of participants providing data for each marker. A participant who deteriorates on any marker is categorised as deteriorating in the multi-marker trajectory classification, regardless of other markers' direction.

Economic estimate is one-sided. Worsening contributes zero dollars (the credit formula floors at zero), but this does not make the worsening invisible. The \$0 contribution and the worsening movement count are both reported. The Forecast is never summarised as "no change" when worsening has been measured.

Non-monetised markers. SBP, glucose, and waist contribute zero dollars for all participants under v1.1 — but their movement direction is still surfaced. A company whose waist or BP markers are collectively worsening will see that in the Workforce-Risk Intelligence output, independent of the dollar figure.

Small-cell suppression. Exact counts below 5 are suppressed (displayed as "<5") to prevent identification of individual employees. Direction classifications (improving/deteriorating/stable) are preserved even under suppression when the cell count ≥ 1 .

§10 HbA1c conditional module (disabled)

The coefficient register includes an HbA1c module: **\$545 per 1.0 percentage-point reduction in HbA1c per person-year**, sourced from Lage & Boye (2020, *Current Medical Research and Opinion*), applicable to participants with diagnosed Type 2 diabetes and a baseline HbA1c $\geq 7.0\%$, with paired valid readings.

This module is **disabled** in v1.1. It does not contribute to any Forecast. It is present in the register to document the module specification and evidence basis so it can be reviewed by an actuary prior to enabling.

Enabling conditions: the HbA1c module may not be enabled without (a) a methodology version bump (v1.2 or later), (b) independent actuarial acceptance of the coefficient and eligibility rules, and (c) a separate T2D diagnosis data source (the Metra platform does not currently store clinical diagnoses — enabling this module would require a new data intake). The registry validation layer enforces that exactly one active economic pathway exists at any time; the HbA1c module is also mutually exclusive with the BMI module per participant.

§11 Interpretation and boundary language

The following interpretation string is required on every Workforce Exposure Forecast output (it is stored verbatim in the coefficient register and reproduced in every generated document):

"Annual associated healthcare-expenditure difference implied by qualifying measured BMI movement; not guaranteed claims savings, premium, carrier pricing, or individual claims prediction."

Every generated Forecast also carries a PROPOSED-status disclosure when the active coefficient is in PROPOSED status:

"The BMI economic pathway coefficient (\$93/BMI-point/person-year) is a proposed parameter subject to independent actuarial review. It has not been independently accepted as of this document's generation date."

§12 Limitations

- L1 ESI proxy.** The Metra platform does not store an employer-sponsored-insurance flag at the employee level. All enrolled employees are treated as ESI (ASSUME_ESI_FOR_ENROLLED = true). Employees who are enrolled but not on ESI (e.g., spouse coverage, public exchange coverage) will be incorrectly included. The size of this error is not quantified.

- L2 Observational association, not causal.** The \$93 coefficient is a per-BMI-point spending association from observational data (Thorpe & Joski 2024). The methodology adopts the conservative lower 95% CI bound, but cannot establish causation between measured BMI movement and future claim reduction. Confounding (selection into weight loss, concurrent interventions, regression to the mean) is not modelled.
- L3 Coefficient uncertainty not propagated.** The bootstrap CI covers only cohort-sampling uncertainty. The \$93 value is used as a point estimate; the CI around it from Thorpe & Joski (2024) is not propagated through the bootstrap. A full propagation would produce wider intervals.
- L4 No nonparticipant extrapolation.** The Forecast is bounded to qualifying measured participants. It does not project what the broader enrolled population might achieve. The participation rate and representation denominator are reported so the reader may contextualise the scope.
- L5 No attrition or trajectory projection.** The point estimate is annualised by coefficient units but does not model future attrition, engagement decay, or trajectory continuation. It represents the exposure implied by the BMI movement measured to date, not a multi-year projection.
- L6 BMI range capped at 45.** Participants with baseline BMI above 45 are outside the evidence boundary of the Thorpe & Joski (2024) association and are excluded. Their exclusion is reported in the denominator table.
- L7 Marker-movement thresholds.** The risk-intelligence layer classifies movement as improving, stable, or deteriorating based on configurable platform thresholds (weight ± 0.5 kg, SBP ± 3 mmHg, glucose ± 5 mg/dL, waist ± 2 cm). These are chosen at measurement-noise scale and are not sourced from the v1.1 evidence documents. The economic estimate is unaffected (it uses absolute BMI values, not threshold buckets).
- L8 HbA1c module disabled.** The T2D HbA1c subgroup module is not active. Employers with a high diagnosed-T2D population receive no dollar signal from HbA1c improvement under v1.1. Enabling requires a separate data intake, a methodology version bump, and actuarial acceptance.

§13 Attestation protocol

Metra publishes this methodology in full — including the coefficient register, eligibility rules, credit formula, bootstrap procedure, boundary language, and all named limitations — so that a credentialed actuary can review it independently and, if they concur, attest to it.

What an attestation covers. A methodology-level attestation is a credentialed actuary's statement that they have reviewed the v1.1 specification and find the BMI economic pathway (\$93/BMI-point, Thorpe & Joski 2024 lower 95% CI bound), the eligibility rules, the credit formula, the deterministic bootstrap procedure, and the named limitations to be reasonable and appropriate for producing a buyer-side forward estimate of workforce BMI-related economic exposure, and that the boundary interpretation language conforms to ASOP Nos. 25, 41, and 56.

Note on the PROPOSED status. Because the BMI coefficient is currently in PROPOSED — PENDING ACTUARIAL REVIEW status, the most impactful attestation a reviewer can provide is an acceptance of the \$93 coefficient and the evidence basis. Upon acceptance by a credentialed actuary, the coefficient status will move to ACTIVE, the PROPOSED disclosure will be removed from generated Forecasts, and this document will be updated.

How to submit. Credentialed actuaries (ASA, FSA, MAAA, FCAS, ACAS) may submit a review or attestation by reading this document in full and contacting Metra at care@usemetra.com with subject line "*Methodology Attestation — v1.1*" or "*Methodology Review — v1.1*". Submissions carry the reviewer's name, credential, firm, credential number, optional firm letterhead, and a statement of scope. Compensation arrangements, where any exist, are disclosed publicly on the [attestation roster](#). Critiques and counter-derivations are equally welcomed and inform v1.2.

Version-specific attestations. A v1.1 attestation is pinned to v1.1. It does not transfer to v1.2 or later. Reviewers who attested to v1.0 are invited to re-attest to v1.1 — the methodology has changed materially and a new review is required.

§14 Version history

Version	Register	Date	Summary
v1.0	original	Prior	Four-marker additive stacking — weight \$68/lb, SBP \$150/mmHg, glucose \$50/mg-dL, waist \$30/cm. Bühlmann credibility blend ($Z = N/(N+400)$ PQ), lag haircut $\lambda = 0.50$, attrition $\alpha = 0.15$, bootstrap $B = 10,000$. Now pinned historical at /methodology/v1.0/ .
v1.1-draft	2026-08-refresh	2026-08	BMI economic pathway only (\$93/BMI-point, Thorpe & Joski 2024 lower 95% CI, PROPOSED). SBP/glucose/waist non-monetised; legacy coefficients retired. No credibility blend in exposure engine. No lag haircut. No horizon projection. Deterministic digest-seeded bootstrap. Versioned register + status lifecycle. Eligibility rules codified. Bidirectional visibility enforced. HbA1c module present but disabled.

§15 Bibliography

1. Thorpe, K.E. & Joski, P. (2024). "Per-BMI-Point Healthcare Spending Association for Employer-Sponsored Insurance Adults." *JAMA Network Open*. [BMI pathway, \$6; \$93/BMI-point lower 95% CI bound, \$5.]
2. Lage, M.J. & Boye, K.S. (2020). "The relationship between HbA1c reduction and healthcare costs among patients with type 2 diabetes." *Current Medical Research and Opinion* 36(6): 1–9. [HbA1c conditional module, \$10; \$545/ppt, \$5.]
3. Kumar, N. et al. (2024). "Categorical hypertension cost burden in employer-sponsored populations." *Hypertension*. [SBP non-monetisation rationale, \$5.]
4. Parker, E. et al. (2024). "Fasting glucose and healthcare expenditure: subgroup and categorical analysis." *Diabetes Care*. [Fasting glucose non-monetisation rationale, \$5.]
5. Ross, R. et al. (2020). "Waist circumference as a vital sign in clinical practice: A Consensus Statement from the IAS and ICCR Working Group on Visceral Obesity."

Nature Reviews Endocrinology 16: 177–189. [Waist circumference non-monetisation rationale, §5.]

6. Boudreau, D.M. et al. (2009). "Cardiovascular-related costs in treated patients with abdominal obesity and hypertension in a managed care setting." *Journal of Managed Care Pharmacy*. [Waist circumference non-monetisation, §5.]
7. Bühlmann, H. (1967). "Experience Rating and Credibility." *ASTIN Bulletin*. [Bühlmann credibility — applied in v1.0; not applied in v1.1 exposure engine. Applied in v2.1 Experience Calibration addendum.]
8. Society of Actuaries. "Credibility Procedures." Educational note. [Credibility — v2.1 only.]

For credentialed actuaries. If you hold an ASA, FSA, MAAA, FCAS, or ACAS credential and are reviewing this document, you are the intended primary audience. The BMI coefficient (\$93/BMI-point/person-year) is in PROPOSED status and the most valuable contribution a reviewer can make is an acceptance or critique of that coefficient and its evidence basis. **Correspondence:** care@usemetra.com, subject line "*Methodology Attestation — v1.1*" or "*Methodology Review — v1.1*".