

Model Card: PriceWise v2.3

Vesta Mutual Insurance AS · Term life and health risk-pricing model · Card version 1.0, 29 June 2026 · Model owner: Chief Actuary · Card maintained by: AI Governance Lead

Worked example for portfolio and training purposes. Vesta Mutual Insurance AS is a fictional company; all data, metrics and names are invented. Prepared by Erik Bernath, Furioso AI Consulting OÜ (furiosoaiconsulting.eu), June 2026. Licensed CC BY 4.0: reuse freely with attribution. This document is informational and is not legal advice.

1. Model details

PriceWise v2.3 is an internally developed risk-scoring model supporting underwriting and pricing of term life and supplementary health policies in Estonia, Latvia and Lithuania. Architecture: gradient-boosted decision trees producing a risk score, benchmarked against a GLM baseline retained for challenge and fallback. Trained January 2026 on portfolio data through September 2025; deployed 1 March 2026. Inference runs in the policy administration system; underwriters see the score, the price band, and the top contributing factors per case.

Regulatory status. High-risk under EU AI Act Annex III 5(c) (risk assessment and pricing of natural persons in life and health insurance). Vesta is both provider (in-house development) and deployer. The provider compliance plan (Annex IV technical documentation, quality management, conformity assessment, EU database registration) is tracked separately; this card is a summary artifact, not the technical documentation itself. The Art. 27 fundamental rights impact assessment for this use is on file. National insurance supervision (Finantsinspektsioon) expectations on actuarial governance apply in parallel.

2. Intended use

- Decision support for underwriters pricing term life and supplementary health applications from natural persons resident in the three Baltic markets.
- Score informs the price band; the underwriter confirms or overrides. Sums assured above €250,000 and all declined cases require 4-eyes review.

3. Out-of-scope uses

- Fully automated issuance, pricing or refusal with no underwriter in the loop.
- Claims decisions of any kind, non-life products, group schemes, or applicants outside EE/LV/LT.
- Any use of score outputs for marketing segmentation or retention pricing.

4. Training data

Pseudonymized portfolio data 2014-2025: approximately 180,000 policies with outcome history. Features: age band, smoker status, BMI band, occupation class, sum assured, disclosed medical conditions (grouped), policy term, and county-level residence. Explicitly excluded as inputs: gender (EU Gender Directive, following the Test-Achats ruling), nationality, language preference, and any geography finer than county. Known data gaps recorded: applicants aged 18-25 are under-represented (4% of training rows), and rural Latvian counties are thin; both have widened manual-review bands as a compensating control.

5. Performance

Segment	Gini coefficient	Calibration ratio	Note
Portfolio overall (holdout, n=36,400)	0.41	1.01	Meets the 0.38 acceptance threshold set by the actuarial function
Age 18-30	0.34	0.93	Below-threshold discrimination and slight under-pricing; thin data; manual review band widened for this segment
Age 31-50	0.43	1.01	Strongest segment
Age 51-65	0.40	1.04	Acceptable; slight over-pricing monitored
Smoker subgroup	0.39	0.99	Acceptable

Rural-county residents	0.38	1.05	Borderline; flagged for next retraining cycle
------------------------	------	------	-----------------------------------------------

Metrics explained. Gini measures the model's ability to separate higher-risk from lower-risk applicants (higher is better). The calibration ratio compares predicted to observed outcomes (1.00 is ideal; above 1 over-prices, below 1 under-prices). Figures are from the 20% temporal holdout; all values in this worked example are invented but shaped to be realistic, including the imperfections.

6. Fairness evaluation

Quarterly tests, run by the actuarial function with results to the AI Governance Committee: (1) proxy leakage: gender is not an input, so the test checks whether occupation class and disclosed conditions reconstruct it; reconstruction AUC of 0.64 in the last test sits under the 0.70 action threshold and is trending stable; (2) price-band distribution across age bands and counties against loss-ratio justification: differences must be explainable by observed risk, and the rural-county over-pricing flag in section 5 came from this test; (3) override-pattern review: whether underwriter overrides concentrate in any demographic, which would indicate the model and human disagree systematically about a group. No red findings to date; the rural calibration item is amber and assigned.

7. Limitations

- The model prices observed history; it can encode historical underwriting conservatism toward groups with thin data, which is why thin segments carry wider manual bands rather than tighter automation.
- Stability: scores for borderline medical-disclosure combinations are sensitive to grouping choices; the grouping table is versioned and any change triggers revalidation.
- Drift: COVID-era mortality and morbidity patterns in the training window are monitored as a known distortion; the 2027 retrain drops pre-2016 rows and re-weights 2020-2022.

8. Human oversight and monitoring

Underwriters see contributing factors with every score and can override with a recorded reason; override rate (currently 7%) is reviewed monthly, and a rate above 15% in any segment pauses the model for that segment pending review. Monitoring: monthly drift dashboard (input distributions, score distributions, calibration on rolling actuals), quarterly fairness tests as above, annual full revalidation by the actuarial function with board reporting. Retraining is scheduled annually or on a red drift or fairness finding, and every retrained version receives a new card before deployment.

9. Card maintenance

This card follows the structure proposed by Mitchell et al. (Model Cards for Model Reporting, 2019), adapted for an insurance deployer-provider under the EU AI Act. It is reviewed at every model version change and at least annually. Version 1.0 covers PriceWise v2.3 only; scores and metrics from other versions are not comparable.