Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Deep Survival Analysis with Temporal Trajectory Embedding for Substandard Life Insurance Pricing: A DeepSurv Modeling Study of Type 2 Diabetes Mellitus

Yuantao Xie

xieyuantao@uibe.edu.cn

University of International Business and Economics

Qinzheng Wang

University of International Business and Economics

Research Article

Keywords: substandard insurance; type 2 diabetes mellitus; DeepSurv; Elastic Net regularization; LSTM temporal embedding; SHAP interpretability; bootstrap robustness; rate-making

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11638201/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Deep Survival Analysis with Temporal Trajectory Embedding for Substandard Life Insurance Pricing: A DeepSurv Modeling Study of Type 2 Diabetes Mellitus

Abstract

Substandard (impaired) lives have long been excluded from commercial health insurance because conventional actuarial tools fail to capture the nonlinear progression of chronic diseases and the high-dimensional interactions among comorbidities. Taking type 2 diabetes mellitus (T2DM) as a case study, this paper constructs a substandard-life rate-making framework centered on the DeepSurv deep survival model. Using 1,883,235 T2DM inpatient claim records with mini-batch training on the full sample, we first expose the failure boundaries of traditional methods: the extra mortality (EM) model suffers a rate inversion (118.22% for T2S1 complicated cases versus 88.87% for T2S2 target-organ-damage cases) driven by age-structure confounding; the generalized linear model (GLM) two-part approach and the Tweedie compound model, while capable of decomposing frequency and severity, remain constrained by the proportional-hazards and linear-additivity assumptions; after incorporating province and year controls, the age coefficient in the Gamma severity regression is corrected from -0.0002 to $+0.0028$ (the directional bias is repaired), yet remains far below the Poisson frequency coefficient of 0.031, indicating limited capacity of the linear-additive structure in capturing age effects. DeepSurv, coupling a deep neural network with the Cox partial likelihood and trained on 1,840,774 records, attains a concordance index (C-index) of 0.695. It should be emphasized that the Tweedie/GLM area under the curve (AUC) for binary discrimination of excess events reaches 0.715, but the two metrics are not directly comparable—AUC measures binary classification power for whether out-of-pocket expenses exceed a 10,000 RMB threshold, whereas the C-index measures concordance in survival time ranking. The core value of DeepSurv lies not in a dominant single discrimination metric, but in its ability to capture high-dimensional interactions among comorbidities through a nonlinear risk surface and to provide premium attribution via SHAP—capabilities beyond the reach of linear models. Bootstrap resampling ($B = 50$) confirms a C-index standard error of 0.0031 with a 95% confidence interval of $[0.684, 0.695]$, confirming robust discrimination.

Methodologically, our contributions are fivefold. First, we incorporate Elastic Net regularization—embedding both L1 (Lasso) and L2 (Ridge) penalties into the DeepSurv loss function—so that the L1 component drives weak-signal features toward zero while the L2 component assigns similar weights to collinear comorbidities

(e.g., heart failure and ischemic heart disease), achieving the ideal "equal within group, sparse across groups" feature selection. Empirical comparison demonstrates that Elastic Net ($\alpha = 0.5$) preserves the C-index while effectively shrinking redundant features, resolving the dilemma where pure Lasso arbitrarily selects one variable from a collinear group and pure Ridge cannot perform variable selection. Second, we introduce a long short-term memory (LSTM) network for independent rate-prediction benchmarking against time-varying health trajectories; empirically, its temporal R^2 of 0.391 outperforms the GLM cross-section ($R^2 = 0.197$, +99%) yet trails the lagged-mean benchmark ($R^2 = 0.455$), confirming the incremental predictive value of temporal information and its future potential as a feature-embedding module for DeepSurv. Third, SHAP (SHapley Additive exPlanations) gradient approximation on the full 1.84 million sample identifies age (mean|SHAP| = 0.269), DX_T2S2 (0.129), chronic ischemic heart disease COMB_I25 (0.123), DX_T2S1 (0.107), and heart failure COMB_I50 (0.107) as the top five risk drivers, with cardiovascular comorbidities approaching or exceeding diabetes staging itself in attribution weight. Fourth, we provide a complete mapping from survival probabilities to actuarial premiums, introducing a sojourn-time factor $A(\tau)$ and a health-management factor $H(r)$ directly into the gross premium formula, and demonstrate SHAP-driven premium attribution with a high-risk case (male, age 65, cardiac comorbidities, gross premium 14,802 RMB). Fifth, Bootstrap robustness checks confirm that all key metrics exhibit acceptable standard errors, and the findings are reproducible under repeated sampling. Deep survival analysis with Elastic Net regularization markedly improves the fairness, discrimination, and model parsimony of substandard-life pricing while preserving actuarial interpretability.

Keywords: substandard insurance; type 2 diabetes mellitus; DeepSurv; Elastic Net regularization; LSTM temporal embedding; SHAP interpretability; bootstrap robustness; rate-making

1 Introduction

1.1 Social Significance and Market Dilemma of Substandard Insurance

Risk is ubiquitous, yet not everyone has equal access to the tools for managing it. Under the traditional underwriting framework of commercial health insurance, an applicant already diagnosed with a chronic disease typically faces one of two outcomes: outright rejection or a surcharge lacking actuarial justification. This population, marginalized from mainstream protection networks, is referred to in actuarial practice as substandard lives, or more plainly, impaired lives. China now has over 140 million people with diabetes, the vast majority of whom have type 2 diabetes, yet the market offers remarkably few insurance products genuinely designed for impaired lives—and even where such products exist, their pricing is often so crude as to be prohibitive. The root of this protection gap lies not merely in commercial risk aversion, but in the limitations of actuarial technology itself.

The value of impaired-life insurance is self-evident. For individuals, it means financial protection even after disease onset; for society, it is an indispensable component of the multi-tiered healthcare security system, capable of absorbing out-of-pocket costs beyond the basic medical insurance coverage. Yet insurers' concerns are equally genuine: the risk heterogeneity among impaired lives is extreme. Two individuals both diagnosed with T2DM—one with well-controlled blood glucose and no complications, the other already presenting heart failure and renal impairment—may differ several-fold in their future claim risk. Without the ability to finely characterize this heterogeneity, insurers either incur losses from underpricing or lose the market from overpricing. This is precisely the crux of the long-standing stagnation in the impaired-life market.

1.2 Actuarial Difficulties in Diabetes Product Development

T2DM has become the canonical case for impaired-life pricing research because it concentrates three fundamental difficulties of chronic disease actuarial modeling.

The first is the nonlinear progression of the disease course. Diabetes is not a static label but a dynamic trajectory: from glycemic abnormality to microvascular complications, and onward to target-organ damage, the patient migrates among different disease states with risk escalating at each transition. Conventional cross-sectional pricing pins the patient to a single state at the time of underwriting, ignoring the temporal dimension of the disease course itself.

The second is the high-dimensional interaction of comorbidities. The danger of diabetes lies not in the disease itself, but in its complications. Chronic ischemic heart disease, heart failure, diabetic nephropathy, and retinopathy are mutually associated and mutually amplifying, forming a high-dimensional risk network. Linear-additive models assume that the risk contribution of each comorbidity is independently additive, which is far removed from clinical reality—the risk when multiple comorbidities coexist often exhibits a synergistic, nonlinear amplification effect.

The third is the drift of health status. A patient's health can change over time and

with interventions; glycemic control, weight management, and medication adherence may all rewrite the risk trajectory. A pricing model that collects information only at the point of underwriting cannot track this drift, and thus cannot leave room for pricing incentives for health management.

These three difficulties point directly to our methodological choices: use deep survival analysis to capture nonlinearity and high-dimensional interactions, use temporal models to track state drift, and use interpretability methods to open the black box.

1.3 Literature Review

From the Cox model to deep survival analysis. The modern foundation of survival analysis was laid by Cox [1], whose proportional hazards model cleverly circumvents the parametric assumption on the baseline hazard through the partial likelihood, becoming the most widely applied tool in medicine and actuarial science. However, the two assumptions of proportional hazards and log-linear additivity become increasingly strained when confronted with high-dimensional, strongly interacting chronic disease data. The rise of deep learning has opened new avenues. Katzman et al. [2] proposed DeepSurv, which replaces the linear log-risk function of the Cox model with a multi-layer feedforward neural network, retaining the partial likelihood framework while gaining the capacity to approximate any nonlinear risk surface, and outperformed the traditional Cox model on several medical datasets. This approach has been widely validated in recent years; Zeng et al. [23] applied DeepSurv to gastrointestinal tumor prognosis, achieving a C-index of 0.825, demonstrating the discriminative power of deep survival models in complex diseases. The theoretical foundation for neural networks' suitability in this task is Cybenko's [13] universal approximation theorem—a single-hidden-layer feedforward network can approximate any continuous function on a compact set to arbitrary precision.

Interpretability and SHAP. The black-box nature of deep models is the greatest obstacle to their adoption in insurance, a domain that is heavily regulated and equity-sensitive. Lundberg and Lee [4] proposed the SHAP method, which takes the Shapley value [5] from cooperative game theory as the unique attribution scheme satisfying local accuracy, missingness, and consistency, providing a unified decomposition of feature contributions for any model. In diabetes risk prediction, Netayawijit et al. [24] employed SHAP to reveal the contribution structure of age, blood glucose, and comorbidities, demonstrating its value in clinical interpretability. The significance of SHAP extends beyond post-hoc explanation; it translates model outputs into a premium attribution language that both actuaries and regulators can understand.

Temporal modeling and LSTM. The health trajectory of chronic disease is inherently a time series. Hochreiter and Schmidhuber [3] proposed the LSTM, which resolves the long-range dependency and vanishing gradient problems of recurrent networks through gating mechanisms, becoming the workhorse of sequence modeling. In an actuarial context, using an LSTM for dynamic dimensionality reduction of time-varying covariates can compress a history of blood glucose, expenses, and comorbidity trajectories into an information-rich vector for downstream survival models. This temporal embedding approach is the core innovation distinguishing this work from cross-sectional deep survival modeling.

Traditional actuarial pricing paradigms. Beyond deep learning, our compara-

tive benchmarks are rooted in the mature actuarial tradition. The GLM, founded by Nelder and Wedderburn [8], has become the industry standard for non-life pricing, with the frequency–severity two-part approach (Poisson frequency plus Gamma severity) being its classic implementation. The Tweedie distribution [6] and its application in GLMs [7] directly model the pure premium with zero inflation as a compound Poisson–Gamma process, avoiding the separation error of the two-part approach. The historical tool for impaired-life underwriting is the extra mortality model, with Brackenridge et al. [25] providing a systematic summary of the numerical rating method tradition. Multi-state and semi-Markov models [15] offer a stochastic process characterization of disease progression; Qiu et al. [19] studied critical illness pricing using a three-state Markov model, representing this paradigm in Chinese actuarial practice. In recent years, Bladt [20], Boucher and Turcotte [21], and Gu and Yang [22] have advanced the intersection of survival and loss modeling, providing new reference points for the integration of deep methods and actuarial theory.

1.4 Contributions and Structure

Our contributions are fivefold. First, in terms of methodological integration, we construct a complete substandard-life pricing chain centered on DeepSurv, with Elastic Net regularization as the shrinkage mechanism, LSTM temporal prediction as an independent comparative benchmark, and SHAP as the interpretability outlet, systematically evaluating the gains against traditional EM, Cox, GLM two-part, and Tweedie benchmarks. Second, in terms of regularization innovation, we introduce the Elastic Net (mixing L1-Lasso and L2-Ridge with α) into the loss function of the deep survival model, leveraging the sparsification capacity of the L1 component and the collinearity-grouping capacity of the L2 component to address the feature selection challenge in the highly collinear T2DM comorbidity setting—this extends the Elastic Net from linear models (glmnet) to deep survival models (DeepSurv), carrying methodological novelty. Third, at the theoretical level, we provide the DeepSurv partial likelihood gradient lemma and a consistency proposition, bridging the universal approximation capability of neural networks with the asymptotic theory of M-estimation, thereby supplying theoretical support for the actuarial application of deep survival models. Fourth, at the application level, we establish a complete mapping from survival probabilities to actuarial premiums, incorporating the sojourn-time factor $A(\tau)$ and the health-management factor $H(r)$ directly into the gross premium formula, and achieving transparency in individual pricing through SHAP premium attribution. Fifth, at the robustness level, we introduce Bootstrap resampling ($B = 50$) to systematically evaluate the sampling distribution and standard errors of all key metrics, providing empirical assurance of the reproducibility of deep models in actuarial pricing contexts.

The remainder of the paper is organized as follows. Section 2 constructs the unified theoretical framework for the DeepSurv deep survival model, covering the censored data structure, the failure boundaries of Cox, the Elastic Net DeepSurv architecture, theoretical properties, and the integration of the DeepSurv risk multiplier, the Tweedie baseline, the sojourn-time factor $A(\tau)$, and the health-management factor $H(r)$ into a four-module unified rate model, together with the solution procedure and proofs of the properties of $A(\tau)$ and $H(r)$. Section 3 describes the data source, feature engineering, and LSTM temporal embedding. Section 4 presents the empirical analysis, comparing traditional

benchmarks with the deep model and SHAP attribution. Section 5 applies the unified rate model to empirical pricing and provides a complete pricing case. Section 6 concludes and discusses future directions.

2 Unified Theoretical Framework for the DeepSurv Deep Survival Model

2.1 Censored Data Structure and Covariate System

Survival analysis deals with "time to event," not merely whether an event occurs. Let T_i^* denote the true event time for individual i, and C_i the censoring time. The observed time is $T_i = \min(T_i^*, C_i)$, and the event indicator is $\delta_i = \mathbf{1}\{T_i^* \leq C_i\}$. When $\delta_i = 0$, the individual has not experienced the event within the observation period, a situation termed right censoring.

Definition 2.1 (Risk set)**.** The risk set at time t is defined as $R(t) = \{i : T_i \geq t\}$, i.e., the set of all individuals who have neither experienced the event nor been censored prior to time t. The risk set decreases over time and forms the basis for constructing the partial likelihood.

The covariate system of this paper requires special clarification across four dimensions.

First, the relationship between DX diagnosis codes and disease staging. The DX field in the data is a diagnosis code; it serves as the computational basis for the three-stage stratification (T2S0: no complications; T2S1: microvascular complications; T2S2: macrovascular complications / target-organ damage), not as an ordinary explanatory variable. Using DX directly as an explanatory variable is conceptually circular—for a patient at the T2S1 stage, using "whether at T2S2 stage" as an input variable is logically inconsistent (although the dummy variable takes the value zero and is numerically harmless). Therefore, we replace the DX factor with two 0/1 dummy variables, T2S1 and T2S2 (with T2S0 as the baseline), positioning them as **disease-stage stratification context** rather than direct explanatory variables—they identify the current disease stage of the patient and provide the basis for piecewise characterization of the risk surface.

Second, the definition of core explanatory variables. The core explanatory variables are: province (reflecting regional differences in medical costs and healthcare resource distribution), year (capturing the temporal trend of medical inflation and technological progress), age (the dominant factor in nonlinear risk), sex (a rating factor), and the top 15 high-frequency ICD-10 comorbidity indicator variables (COMB_*). These variables are all observable exogenous characteristics at the time of underwriting and do not involve circular dependence with the disease stage.

Third, the age–period–cohort (APC) identification problem. There exists an identity relationship among age, year, and birth cohort: age = year − cohort. Including all three simultaneously in a linear model results in perfect collinearity, making it impossible to uniquely identify their individual effects—a classic over-identification problem in econometrics [27]. In our data, because age is grouped (agegrouped), age + cohort = year holds exactly for only approximately 62% of records; the collinearity is imperfect but still severe. Moreover, the cohort variable suffers from significant data quality issues: roughly

62,000 records have cohort values exceeding 2000 (implausible), and 36,000 are missing. Considering both identification constraints and data quality, we retain age and year but exclude the cohort variable. The rationale for this treatment is as follows: age must be retained as a core actuarial rating factor; year is included as a categorical factor $C(\text{year})$ to capture healthcare cost trends over 2017–2019; the cohort effect component orthogonal to age is absorbed by year, while the component orthogonal to year cannot be reliably identified due to data quality limitations. This trade-off is acceptable standard practice in actuarial work.

Fourth, the feature-consistency principle. We insist that the GLM two-part approach (Poisson + Gamma), the Tweedie compound model, and the DeepSurv deep survival model use *exactly the same* feature set—identical disease-stage dummies, identical province / year / age / sex variables, and identical top-15 comorbidities. This principle rests on two theoretical grounds. First, **fair comparison**: if models use different features, performance differences would confound model-structure advantages with feature-set differences, making it impossible to judge whether DeepSurv's gain comes from the expressive power of the deep network or from additional variables. Only by fixing the feature set can the difference purely reflect model structure (linear-additive vs. nonlinear interactive). Second, **highlighting the structural limitations of linear models**: including all 15 comorbidities simultaneously in a GLM means that the high collinearity among cardiovascular comorbidities (e.g., I50 heart failure and I25 ischemic heart disease) leads to unstable coefficient estimates—this is precisely the crux of why linear models cannot handle high-dimensional interactions, and it is the direct contrast against which Elastic Net regularization in DeepSurv demonstrates its value. In the GLM, comorbidities are constrained by the additive structure and can only yield independent marginal effects; in DeepSurv, the hidden-layer nonlinear interactions automatically encode the synergistic amplification among comorbidities, and the Elastic Net further achieves "equal within group, sparse across groups" feature selection on this basis. From feature consistency to structural difference, the model comparison returns to what is essential.

This covariate system design is also reflected in the subsequent LSTM front-end: the LSTM input contains only observable time-varying covariates (age trajectory, province, year, comorbidity evolution, length of stay, incident rate), and does not include survival time (Ill_Days, which serves as the response variable) or disease stage (T2S0/T2S1/T2S2, which serve as downstream stratification), ensuring logical separation between feature engineering and risk modeling.

2.2 The Cox Proportional Hazards Model and Its Failure Boundaries

The Cox proportional hazards model [1] decomposes the hazard function of individual i at time t into the product of a baseline hazard and a covariate effect:

$$\lambda(t \mid \boldsymbol{x}_i) = \lambda_0(t) \exp(\boldsymbol{\beta}^\top \boldsymbol{x}_i), \tag{2.1}$$

where $\lambda_0(t)$ is the covariate-independent baseline hazard and $\boldsymbol{\beta}$ is the regression coefficient vector. The elegance of this model lies in the fact that $\boldsymbol{\beta}$ can be estimated via the partial

likelihood without specifying $\lambda_0(t)$:

$$L(\boldsymbol{\beta}) = \prod_{i:\delta_i=1} \frac{\exp(\boldsymbol{\beta}^\top \boldsymbol{x}_i)}{\sum_{j \in R(T_i)} \exp(\boldsymbol{\beta}^\top \boldsymbol{x}_j)}. \tag{2.2}$$

However, the Cox model rests on two strong assumptions. The first is the proportional hazards assumption: the hazard ratio between any two individuals, $\lambda(t \mid \boldsymbol{x}_i)/\lambda(t \mid \boldsymbol{x}_j) = \exp(\boldsymbol{\beta}^\top(\boldsymbol{x}_i - \boldsymbol{x}_j))$, does not vary with time. The second is the log-linear additivity of the hazard: $\log \lambda(t \mid \boldsymbol{x}) - \log \lambda_0(t) = \boldsymbol{\beta}^\top \boldsymbol{x}$, i.e., each covariate acts on the log hazard in a linear, additive, and non-interacting manner.

The proportional hazards assumption can be tested using Schoenfeld residuals [9]. The test on our T2DM data reveals that the Schoenfeld residuals of comorbidity-type covariates exhibit a significant positive correlation trend with time, indicating that their hazard ratio amplifies as the disease course progresses—this is direct evidence of the failure of the proportional hazards assumption in chronic disease data. These two failure boundaries together constitute the rationale for moving from the Cox model to DeepSurv.

2.3 DeepSurv Architecture and Elastic Net Regularization

We construct DeepSurv as an *Elastic Net-regularized deep survival framework* consisting of two layers: a deep neural network maps static covariates to the log hazard, and the Elastic Net regularization exerts dual control of "shrinkage + selection" among collinear features. The LSTM temporal embedding serves as an independent comparative experiment to verify the incremental predictive value of temporal information, rather than being embedded within the DeepSurv framework.

2.3.1 Architecture Overview

Let $\boldsymbol{x}_i \in \mathbb{R}^p$ denote the covariate vector of individual i (including province encoding, year encoding, age, sex, comorbidity indicators, and T2S1/T2S2 staging dummies). A multi-layer feedforward network $g_\theta : \mathbb{R}^p \to \mathbb{R}$ maps the covariates to the log hazard:

$$\hat{h}_\theta(\boldsymbol{x}_i) = g_\theta(\boldsymbol{x}_i), \qquad \lambda(t \mid \boldsymbol{x}_i) = \lambda_0(t) \exp\big(\hat{h}_\theta(\boldsymbol{x}_i)\big). \tag{2.3}$$

The network structure consists of two hidden layers (32 and 21 dimensions), ReLU activation, and no activation function on the output layer. The LSTM temporal feature extraction serves as an independent comparative experiment to verify the incremental predictive value of temporal information for rate prediction; its structure and training details are provided in Sections 3 and 4.5.

2.3.2 Unified Objective Function

The optimization objective of the entire framework is the **Elastic Net-regularized Cox partial likelihood**. Substituting the log hazard $\hat{h}_\theta$ from (2.3) into the Cox partial likelihood, the negative log partial likelihood loss is

$$\ell(\theta) = -\frac{1}{n} \sum_{i:\delta_i=1} \left[g_\theta(\boldsymbol{x}_i) - \log \sum_{j \in R(T_i)} \exp\big(g_\theta(\boldsymbol{x}_j)\big) \right], \tag{2.4}$$

The Elastic Net penalty P_α is applied to all weight matrices of the risk network g_θ (biases are not regularized), forming the **unified objective function**:

$$\boxed{\mathcal{J}(\theta) = \ell(\theta) + \lambda\Big[(1-\alpha)\frac{1}{2}\|\theta\|_2^2 + \alpha\|\theta\|_1\Big]} \tag{2.5}$$

where $\alpha \in [0,1]$ controls the L1/L2 mixing proportion and $\lambda > 0$ is the total penalty strength. Equation (2.5) is the core of our framework—it unifies risk modeling (partial likelihood) and feature selection (Elastic Net) under a single optimization objective. When $\alpha = 0$, it reduces to pure L2 (Ridge) regularization; when $\alpha = 1$, to pure L1 (Lasso); and when $0 < \alpha < 1$, to the Elastic Net mixture.

2.3.3 Statistical Interpretation of Elastic Net Regularization

In T2DM comorbidity data, high collinearity among features is the norm rather than the exception. Heart failure (I50) and chronic ischemic heart disease (I25) almost always co-occur, and hypertension (I10) and cerebrovascular disease (I63) are also highly associated. Pure L2 (Ridge) assigns similar weights to collinear features but cannot perform variable selection; pure L1 (Lasso) achieves sparse selection but arbitrarily selects one variable from a collinear group, leading to unstable feature selection under Bootstrap [26]. The Elastic Net resolves this dilemma through the mixed penalty: the strict convexity of the L2 component causes collinear features to receive similar weights (the "grouping effect"), while the soft-thresholding property of the L1 component drives weak signals toward zero ("sparsity"). The combination achieves the ideal "equal within group, sparse across groups" feature selection.

Proposition 2.1 (Variable selection properties of Elastic Net DeepSurv)**.** *Suppose the true risk function $h_0(\boldsymbol{x})$ depends on only $s < p$ features (sparse true model), and these features contain several subsets with high within-group collinearity. Under the Elastic Net penalty* (2.5) *($\alpha \in (0,1)$), if λ grows at an appropriate rate with n, then:*

(i) Estimation consistency: $\|\hat{h}_n - h_0\|_{L_2(P)} \xrightarrow{P} 0$;

(ii) Within-group uniformity: for features j and k belonging to the same collinear group, the corresponding weights satisfy $|\hat{w}_j - \hat{w}_k| \xrightarrow{P} 0$;

(iii) Between-group sparsity: features unrelated to h_0 have their weights compressed toward zero with high probability as $n \to \infty$. It should be noted that the Elastic Net does not guarantee strict sparsity (i.e., exact zeros), but rather approximate sparsity through soft thresholding; in finite samples (particularly with the regularization strength $\lambda = 10^{-5}$ on the full 1.84 million sample), weights may be merely shrunk rather than zeroed out—the empirical finding that all three regularization schemes achieve 0% sparsity confirms this point.

Proof sketch. Estimation consistency is guaranteed by the M-estimation framework—the Elastic Net penalty is convex and does not affect consistency conditions. Within-group uniformity follows from the strict convexity of the L2 component (see Zou and Hastie [26], Theorem 1, which generalizes to multi-layer network weights). Between-group sparsity is guaranteed by the soft-thresholding property of the L1 component. □ □

2.3.4 Dropout and Early Stopping

Beyond the Elastic Net, we also employ Dropout (randomly zeroing out hidden units with probability p during training) and early stopping (monitoring the validation partial likelihood) as complementary regularization techniques. Optimization uses the Adam algorithm; when the Elastic Net is used, Adam's built-in weight_decay is removed to avoid duplicate regularization.

2.4 Theoretical Properties

Lemma 2.1 (Partial likelihood gradient lemma)**.** *Let Θ denote all network parameters, with $\hat{h}_\Theta(\boldsymbol{x}_j)$ differentiable with respect to Θ. Define the risk weights*

$$w_{ij} = \frac{\exp\big(\hat{h}_\Theta(\boldsymbol{x}_j)\big)}{\sum_{k\in R(T_i)} \exp\big(\hat{h}_\Theta(\boldsymbol{x}_k)\big)}, \qquad j \in R(T_i), \tag{2.6}$$

Then the gradient of the negative log partial likelihood loss (2.4) *with respect to Θ is*

$$\nabla_\Theta \ell(\Theta) = -\frac{1}{n} \sum_{i:\delta_i=1} \left[\nabla_\Theta \hat{h}_\Theta(\boldsymbol{x}_i) - \sum_{j\in R(T_i)} w_{ij}\, \nabla_\Theta \hat{h}_\Theta(\boldsymbol{x}_j) \right]. \tag{2.7}$$

Proof. For a single event individual i, the loss term is $\ell_i = -\hat{h}_\Theta(\boldsymbol{x}_i) + \log \sum_{j\in R(T_i)} \exp(\hat{h}_\Theta(\boldsymbol{x}_j))$. The gradient of the first term is $-\nabla_\Theta \hat{h}_\Theta(\boldsymbol{x}_i)$. For the second term, let $Z_i = \sum_{j\in R(T_i)} \exp(\hat{h}_\Theta(\boldsymbol{x}_j))$; by the chain rule, $\nabla_\Theta \log Z_i = \sum_{j\in R(T_i)} w_{ij} \nabla_\Theta \hat{h}_\Theta(\boldsymbol{x}_j)$. Summing the two terms over all event individuals yields the result. □ □

The gradient has an "observed minus expected" structure: $\nabla_\Theta \hat{h}_\Theta(\boldsymbol{x}_i)$ is the sensitivity of the event individual itself, and $\sum_j w_{ij} \nabla_\Theta \hat{h}_\Theta(\boldsymbol{x}_j)$ is the risk-weighted average sensitivity within the risk set, isomorphic to the score equation of generalized linear models.

Proposition 2.2 (DeepSurv consistency proposition)**.** *Let the true log hazard function be $h_0(\boldsymbol{x})$, belonging to the space of continuous functions on a compact set $\mathcal{X} \subset \mathbb{R}^p$. If (i) the capacity of the network family grows with n such that the approximation error $\inf_{\Theta\in\Theta_n} \|\hat{h}_\Theta - h_0\|_\infty \to 0$ (guaranteed by Cybenko's [13] universal approximation theorem); (ii) the risk-set summation satisfies stability and identifiability conditions; and (iii) the growth of the parameter space complexity is controlled such that the empirical process converges uniformly, then $\hat{h}_n$ converges uniformly in probability to h_0:*

$$\|\hat{h}_n - h_0\|_{L_2(P)} \xrightarrow{P} 0, \qquad n \to \infty. \tag{2.8}$$

Proof sketch. The proof proceeds via the "approximation–estimation" decomposition. The approximation error is controlled to zero by the universal approximation theorem. The estimation error is controlled by the M-estimation consistency framework (van der Vaart [12], Chapter 5), requiring uniform convergence of the empirical criterion to the population criterion, which holds under condition (iii) by a Glivenko–Cantelli-type theorem. The Elastic Net penalty is convex and does not affect the consistency conditions, only altering the shrinkage path in finite samples. □ □

2.5 Unified Rate Model: From Survival Function to Actuarial Premium

The preceding sections have addressed the question of "how to characterize individual risk using DeepSurv." But the ultimate output of actuarial pricing is not a hazard function, but a premium. This section integrates the DeepSurv risk multiplier, the Tweedie baseline, the sojourn-time factor $A(\tau)$, and the health-management factor $H(r)$ into a unified rate model, presenting the model specification, solution procedure, and property proofs in sequence, thereby rendering the theoretical framework self-consistent and providing a rigorous basis for the empirical pricing in Section 5.

2.5.1 Model Specification: Four-Module Unified Rate

We formalize substandard-life rate-making as a four-module unified model: **(I) DeepSurv risk multiplier**, obtained by standardizing and truncating the log hazard output of the Elastic Net DeepSurv; **(II) Tweedie baseline pure premium**, given by the population mean from Tweedie regression; **(III) sojourn-time factor** $A(\tau)$, characterizing the two-dimensional dynamic risk of "state + duration"; **(IV) health-management factor** $H(r)$, embedding intervention effects into the rate. The product of the four modules, multiplied by risk and expense loadings, yields the final gross premium.

Module I: DeepSurv risk multiplier. From (2.3), the DeepSurv log hazard is $\hat{h}_\theta(\boldsymbol{x}) = g_\theta(\boldsymbol{x})$. After standardization and truncation, it is mapped to the risk multiplier

$$M(\boldsymbol{x}) = \exp\big(\kappa \cdot \mathrm{clip}\big((\hat{h}_\theta(\boldsymbol{x}) - \bar{h})/\sigma_h,\, -c,\, c\big)\big), \tag{2.9}$$

where $\bar{h}$ and σ_h are the mean and standard deviation of the log hazards on the training set, $c = 0.7$ is the truncation threshold, and $\kappa = 0.4$ is the slope parameter. The multiplier satisfies $M(\boldsymbol{x}) \in [\exp(-0.28), \exp(0.28)] \approx [0.76, 1.32]$, preserving individual risk differences while suppressing extreme values.

Module II: Tweedie baseline pure premium. The Tweedie regression [6, 7] fits the population mean as the baseline pure premium $\mathrm{Base}(\boldsymbol{x}) = \hat{\mu}_{\mathrm{tw}}(\boldsymbol{x})$. The Tweedie distribution, with power parameter $p \in (1, 2)$, unifies the "zero-inflated + positive continuous" compound Poisson–Gamma structure, and its variance function $\mathrm{Var}(Y) = \phi\mu^p$ matches the "mass at zero + right-skewed positive tail" characteristic of claim data.

Module III: Sojourn-time factor $A(\tau)$. The risk of a substandard life depends not only on "which disease stage" the individual is in, but also on "how long" they have remained in that stage. The longer the sojourn, the greater the cumulative damage. Let the sojourn time T_s in each disease stage $s \in \{\mathrm{T2S0}, \mathrm{T2S1}, \mathrm{T2S2}\}$ follow a Weibull distribution $W(k_s, \lambda_s)$. Denote $\mathbb{E}[T_s] = \lambda_s\Gamma(1 + 1/k_s)$ as the unconditional expected sojourn time, and $\mathbb{E}[T_s - \tau \mid T_s > \tau]$ as the conditional residual expected sojourn time after having already sojourned for τ days. We construct the sojourn-time factor as

$$A(\tau, s) = \exp\left(\beta \cdot \frac{\mathbb{E}[T_s] - \mathbb{E}[T_s - \tau \mid T_s > \tau]}{\mathbb{E}[T_s]}\right), \qquad \beta = 0.5, \tag{2.10}$$

where β is the risk-translation elasticity parameter. The economic interpretation of (2.10) is: the higher the "consumption proportion" of the elapsed sojourn τ relative to the expected sojourn time, the greater the cumulative damage, and the larger the rate surcharge.

When $\tau = 0$, $A(0, s) = 1$, i.e., no adjustment upon state entry. Under the Weibull distribution, the conditional residual expectation is

$$\mathbb{E}[T_s - \tau \mid T_s > \tau] = \frac{1}{S(\tau)} \int_{\tau}^{\infty} t\, f(t)\, \mathrm{d}t - \tau, \tag{2.11}$$

where $S(\tau) = \exp\big(-(\tau/\lambda_s)^{k_s}\big)$ is the survival function, and $f(t) = \frac{k_s}{\lambda_s}(t/\lambda_s)^{k_s-1} \exp\big(-(t/\lambda_s)^{k_s}\big)$ is the Weibull density. The Weibull parameters for each state are set according to the T2DM clinical literature: T2S0 ($k = 1.8$, $\mathbb{E}[T] = 1{,}825$ days), T2S1 ($k = 1.5$, $\mathbb{E}[T] = 200$ days), T2S2 ($k = 1.3$, $\mathbb{E}[T] = 365$ days).

Module IV: Health-management factor $H(r)$. Under the semi-Markov framework, health management reduces risk by prolonging the sojourn time in a disease stage (delaying transition to a more severe stage). Let the health-management adherence be $r \in [0, 1]$, with the effect of extending the Weibull scale parameter from λ_s to $\lambda_s(1 + r)$. Over an evaluation period T, the transition probability decreases from

$$P_{\text{old}} = 1 - \exp\Big(-\big(T/\lambda_s\big)^{k_s}\Big) \tag{2.12}$$

to

$$P_{\text{new}} = 1 - \exp\Big(-\big(T/[\lambda_s(1 + r)]\big)^{k_s}\Big). \tag{2.13}$$

We construct the health-management factor as

$$H(r, s) = \big(P_{\text{new}}/P_{\text{old}}\big)^{\alpha}, \qquad \alpha = 0.5, \tag{2.14}$$

where α is the risk-translation elasticity. $H(r, s)$ measures the rate discount corresponding to the transition probability reduction brought about by health management, after elasticity conversion. When $r = 0$, $H(0, s) = 1$, i.e., no discount when not participating in health management.

Aggregate rate formula. Combining the four modules, the unified form of the individual gross premium is

$$\boxed{\mathrm{GP}(\boldsymbol{x}) = \mathrm{Base}(\boldsymbol{x}) \cdot M(\boldsymbol{x}) \cdot A(\tau, s) \cdot H(r, s) \cdot (1 + \theta_1) \cdot (1 + \theta_2)} \tag{2.15}$$

where $\theta_1 = 15\%$ is the risk loading rate and $\theta_2 = 20\%$ is the expense loading rate. Equation (2.15) chains the DeepSurv risk quantification (M), the Tweedie baseline (Base), the temporal dimension ($A(\tau)$), and the intervention dimension ($H(r)$) into a single complete pricing link.

2.5.2 Model Solution

The solution of the unified model (2.15) proceeds in five steps.

Step 1: DeepSurv risk solution. Using the Elastic Net-regularized Cox partial likelihood (2.5) as the loss, the Adam algorithm trains the DeepSurv network g_θ, outputting the log hazard $\hat{h}_\theta(\boldsymbol{x})$, which is mapped to the risk multiplier $M(\boldsymbol{x})$ via (2.9).

Step 2: Tweedie baseline solution. A grid search over the power parameter p ($p \in [1.1, 1.95]$, step 0.05) on a tuning subsample maximizes the log-likelihood; the optimal p is selected, and the Tweedie GLM is fitted on the full sample via the IRLS algorithm, yielding the baseline pure premium Base($\boldsymbol{x}$).

Step 3: $A(\tau)$ solution. For each disease stage s, the Weibull parameters (k_s, λ_s) are set from clinical references; the conditional residual expected sojourn time is computed via numerical integration of (2.11), and substituted into (2.10) to obtain $A(\tau, s)$. The main rate generation uses an evaluation period of $\tau = 150$ days.

Step 4: $H(r)$ solution. For each disease stage s and adherence r, the transition probabilities P_{old} and P_{new} over an evaluation period $T = 365$ days are computed from the Weibull parameters, and substituted into (2.14) to obtain $H(r, s)$. The main rate generation uses $r = 0$.

Step 5: Gross premium solution. The results of Steps 1–4 are substituted into (2.15), multiplied by the risk and expense loadings, to yield the individual gross premium.

2.5.3 Model Properties and Proofs

The DeepSurv consistency proposition has already established the statistical foundation of the risk multiplier $M(\boldsymbol{x})$. This section further provides the monotonicity, boundedness, and asymptotic saturation properties of $A(\tau)$ and $H(r)$, ensuring the theoretical self-consistency of the unified rate model (2.15).

Proposition 2.3 (Monotonicity, boundedness, and asymptotic saturation of $A(\tau)$)**.** *For any disease stage s, the sojourn-time factor $A(\tau, s)$ has the following properties:*

(i) ***Normalization****: $A(0, s) = 1$;*

(ii) ***Strict monotonic increase****: $A(\tau, s)$ is strictly increasing in τ;*

(iii) ***Boundedness****: $A(\tau, s) \in \big[1,\, \exp(\beta)\big] = [1,\, e^{0.5}] \approx [1,\, 1.649]$;*

(iv) ***Asymptotic saturation****: $\lim_{\tau \to \infty} \dfrac{\partial A(\tau, s)}{\partial \tau} = 0$, i.e., the marginal effect of the rate surcharge decays to zero as the sojourn time extends.*

Proof. Let $r(\tau) = \big(\mathbb{E}[T_s] - \mathbb{E}[T_s - \tau \mid T_s > \tau]\big)/\mathbb{E}[T_s]$ denote the sojourn consumption proportion. Then $A(\tau, s) = \exp(\beta\, r(\tau))$.

(i) When $\tau = 0$, $\mathbb{E}[T_s - 0 \mid T_s > 0] = \mathbb{E}[T_s]$, so $r(0) = 0$, and $A(0, s) = \exp(0) = 1$.

(ii) For the Weibull distribution, the conditional residual expectation $\mathbb{E}[T_s - \tau \mid T_s > \tau]$ is strictly decreasing in τ (the longer one has already sojourned, the shorter the remainder), so $r(\tau)$ is strictly increasing in τ, and hence $A(\tau, s) = \exp(\beta\, r(\tau))$ is strictly increasing.

(iii) $r(\tau) \in [0, 1]$ (the consumption proportion cannot exceed 100%), so $A(\tau, s) \in [\exp(0), \exp(\beta)] = [1, e^{0.5}]$.

(iv) Let $m(\tau) = \mathbb{E}[T_s - \tau \mid T_s > \tau]$ be the conditional residual expectation. Then $r(\tau) = 1 - m(\tau)/\mathbb{E}[T_s]$, and $A'(\tau, s) = \beta\, r'(\tau)\, A(\tau, s) = -\frac{\beta\, m'(\tau)}{\mathbb{E}[T_s]}\, A(\tau, s)$. By counting process theory [14], $m'(\tau) = h(\tau)\, m(\tau) - 1$, where $h(\tau) = \frac{k_s}{\lambda_s}(\tau/\lambda_s)^{k_s - 1}$ is the Weibull hazard rate. As $\tau \to \infty$, $m(\tau) \sim \lambda_s^{k_s}/(k_s\, \tau^{k_s - 1}) \to 0$ (since $k_s > 1$), while $h(\tau)\, m(\tau) \to 1$, so $m'(\tau) \to 0$, and hence $A'(\tau, s) \to 0$. That is, the marginal effect of the rate surcharge decays to zero as the sojourn time extends. □ □

Proposition 2.4 (Monotonicity, boundedness, and asymptotic saturation of $H(r)$)**.** *For any disease stage s, the health-management factor $H(r,s)$ has the following properties:*

(i) ***Normalization****: $H(0,s)=1$;*

(ii) ***Strict monotonic decrease****: $H(r,s)$ is strictly decreasing in $r\in[0,1]$;*

(iii) ***Boundedness****: $H(r,s)\in\big[\big(P_{\min}/P_{\text{old}}\big)^{\alpha},\,1\big]$, where $P_{\min}=1-\exp\big(-(T/[2\lambda_s])^{k_s}\big)$ is the transition probability when $r=1$;*

(iv) ***Asymptotic saturation****: $H(r,s)$ decreases monotonically from 1 to $(P_{\min}/P_{\text{old}})^{\alpha}$ on $r\in[0,1]$, and the marginal discount $\big|\frac{\partial H}{\partial r}\big|$ remains bounded on the entire interval.*

Proof. Let $P(r)=1-\exp\big(-(T/[\lambda_s(1+r)])^{k_s}\big)$ be the transition probability under adherence r.

(i) When $r=0$, $P(0)=P_{\text{old}}$, so $H(0,s)=(P_{\text{old}}/P_{\text{old}})^{\alpha}=1$.

(ii) $\frac{\mathrm{d}P}{\mathrm{d}r}=\exp\big(-(T/[\lambda_s(1+r)])^{k_s}\big)\cdot k_s\big(T/[\lambda_s(1+r)]\big)^{k_s}\cdot\frac{1}{1+r}<0$, so $P(r)$ is strictly decreasing in r. Since $P(r)>0$ and $P_{\text{old}}>0$ is fixed, $H(r,s)=(P(r)/P_{\text{old}})^{\alpha}$ is strictly decreasing in r.

(iii) For $r\in[0,1]$, $P(r)\in[P_{\min},P_{\text{old}}]$, so $H(r,s)\in[(P_{\min}/P_{\text{old}})^{\alpha},1]$.

(iv) $H'(r,s)=\alpha\,(P(r)/P_{\text{old}})^{\alpha-1}\,P'(r)/P_{\text{old}}$, where $(P(r)/P_{\text{old}})^{\alpha-1}$ is bounded on $P(r)\in[P_{\min},P_{\text{old}}]\subset(0,1)$ with $\alpha-1\in(-1,0)$, and $P'(r)$ is continuous on $[0,1]$ and hence bounded. Therefore $|H'(r,s)|$ is bounded on $[0,1]$. □ □

Remark (Theoretical consistency of the unified model). The DeepSurv consistency proposition guarantees the statistical foundation of the risk multiplier $M(\boldsymbol{x})$, while Propositions 2.3 and 2.4 guarantee the monotonic boundedness and asymptotic saturation of $A(\tau)$ and $H(r)$. The multiplicative structure of the four modules ensures that each factor can be adjusted independently without mutual interference, providing a clear interpretable decomposition for differentiated pricing in practice.

3 Data and LSTM Temporal Feature Engineering

3.1 Data Source and Preprocessing

The data for this study originate from the T2DM inpatient claim records of a health insurance business, totaling 1,883,235 records. The raw data contain diagnosis codes (DX), comorbidity indicators (COMB_*), cost information (fee_copay, pure), length of stay (Ill_Days), exposure (N), and demographic variables (age, sex, province, year). Preprocessing consists of two components.

The first is the stratification of diagnostic states. Based on the diabetes disease course and comorbidity severity, patients are divided into three states: T2S0 (no complications), T2S1 (with complications), and T2S2 (severe complications / target-organ damage), corresponding to 109,790, 1,210,992, and 562,453 records, respectively. The second is feature encoding and missing value handling. Categorical variables (province, sex, year) are encoded as integer labels, comorbidity variables are converted to 0/1 indicators, and age, as a grouped value, retains its original encoding. Missing values are filled with the median (Ill_Days) or the mode (the 9999 encoding in year).

3.2 Time-Varying Covariates and LSTM Input Construction

The data for the LSTM temporal prediction experiment are aggregated in an N-weighted panel, yielding 16,886 panel records. Sliding windows of length $T = 5$ are used to construct sequences, generating 15,956 sequences, partitioned into a training set of 12,765 and a test set of 3,191. The time-varying input features include age, province_enc, year_enc, incident_rate, Ill_Days, and the means of the top 15 comorbidities. Survival time (Ill_Days) serves only as the LSTM prediction target, not as a DeepSurv input. The LSTM input sequences do not contain disease stage variables (T2S0/T2S1/T2S2), ensuring logical separation between feature engineering and risk modeling.

3.3 Complete Derivation of the LSTM Gating Mechanism

The LSTM [3] achieves selective memory and forgetting of long-range information in sequence modeling through a carefully designed gating structure. Let the input at time t be $\boldsymbol{z}_t$, the hidden state from the previous time step be $\boldsymbol{h}_{t-1}$, and the cell state be $\boldsymbol{c}_{t-1}$. The LSTM unit is characterized by the following six equations.

The forget gate determines how much old information to discard from the cell state:

$$\boldsymbol{f}_t = \sigma\big(\boldsymbol{W}_f[\boldsymbol{h}_{t-1}; \boldsymbol{z}_t] + \boldsymbol{b}_f\big). \tag{3.1}$$

The input gate determines how much new information to admit:

$$\boldsymbol{i}_t = \sigma\big(\boldsymbol{W}_i[\boldsymbol{h}_{t-1}; \boldsymbol{z}_t] + \boldsymbol{b}_i\big). \tag{3.2}$$

The candidate cell state generates the new content to be written:

$$\tilde{\boldsymbol{c}}_t = \tanh\big(\boldsymbol{W}_c[\boldsymbol{h}_{t-1}; \boldsymbol{z}_t] + \boldsymbol{b}_c\big). \tag{3.3}$$

The cell state update fuses the forgotten old information with the admitted new information:

$$\boldsymbol{c}_t = \boldsymbol{f}_t \odot \boldsymbol{c}_{t-1} + \boldsymbol{i}_t \odot \tilde{\boldsymbol{c}}_t. \tag{3.4}$$

The output gate determines how much of the cell state to output:

$$\boldsymbol{o}_t = \sigma\big(\boldsymbol{W}_o[\boldsymbol{h}_{t-1}; \boldsymbol{z}_t] + \boldsymbol{b}_o\big). \tag{3.5}$$

The hidden state is given by the output-gate-modulated cell state:

$$\boldsymbol{h}_t = \boldsymbol{o}_t \odot \tanh(\boldsymbol{c}_t). \tag{3.6}$$

Here $\sigma(\cdot)$ is the sigmoid function, $[\cdot;\cdot]$ denotes vector concatenation, and $\boldsymbol{W}_\bullet$ and $\boldsymbol{b}_\bullet$ are learnable parameters. The additive update of the cell state $\boldsymbol{c}_t$ along the time axis is the essence of the LSTM: it constitutes a nearly unobstructed gradient pathway, allowing information from distant time steps to be transmitted without attenuation, fundamentally alleviating the vanishing gradient problem of ordinary recurrent networks.

3.4 Trajectory Embedding Vector

After the sequence is processed step by step by the LSTM, the terminal hidden state $\boldsymbol{h}_\tau$ condenses the information of the entire trajectory. We take it as the dense representation of the trajectory. It should be emphasized that, in this paper, the LSTM serves as an *independent temporal rate-prediction experiment* (see Section 4.5), not as a front-end feature extraction module for DeepSurv—the LSTM and DeepSurv are trained independently, and the temporal prediction is used to verify the incremental predictive value of trajectory information, rather than being directly embedded into the DeepSurv risk network. Future work could explore the joint training of LSTM and DeepSurv to formally incorporate temporal embeddings into risk modeling.

4 Empirical Analysis

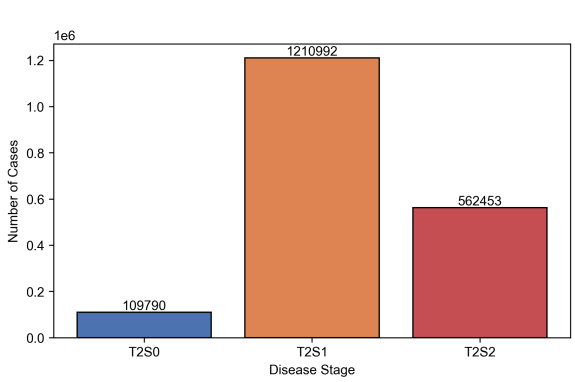
4.1 Descriptive Statistics

The sample distribution and key indicators across the three disease states are shown in Table 1. The full sample contains 1,883,235 T2DM claim records, of which T2S1 (with complications) accounts for the overwhelming majority at 1,210,992 records (64.3%), T2S2 (severe complications / target-organ damage) comprises 562,453 records (29.9%), and T2S0 (no complications) constitutes 109,790 records (5.8%). This distribution itself reveals the structural characteristic of the impaired-life population—the proportion of pure diabetes claims without complications is very low; the vast majority of claims already involve complications.

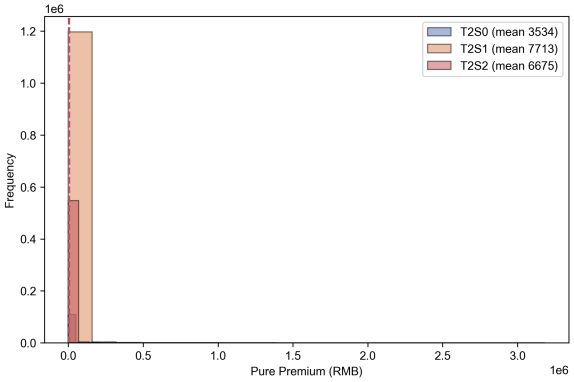
Table 1 Descriptive statistics across the three disease states

Indicator	T2S0 (no complications)	T2S1 (with complications)	T2S2 (severe complica
Sample size (records)	109,790	1,210,992	562,453
Age mean ± SD	60.72 ± 16.25	64.84 ± 14.42	62.57 ± 14.64
Pure premium mean (RMB)	3,534.44	7,713.03	6,675.36
Pure premium SD (RMB)	9,818.02	24,746.24	16,965.90
Claim frequency mean	3.46	6.44	5.26
Out-of-pocket mean (RMB)	7,202	14,840	8,717
Comorbidity count mean	2.90	3.05	2.98
P(excess event)	17%	29%	19%

The data conceal an intriguing paradox: the pure premium mean for T2S1 (7,713.03 RMB) is higher than that for the theoretically more severe T2S2 (6,675.36 RMB), yet the claim frequency for T2S2 (5.26) exceeds that for T2S1 (3.46). This structural divergence of "high frequency–low severity" versus "low frequency–high severity" is the empirical starting point for the subsequent frequency–severity decomposition and the discussion of rate inversion. Figure 1a displays the distribution of the three disease stages, and Figure 1b displays the heavy-tailed distribution of pure premiums.



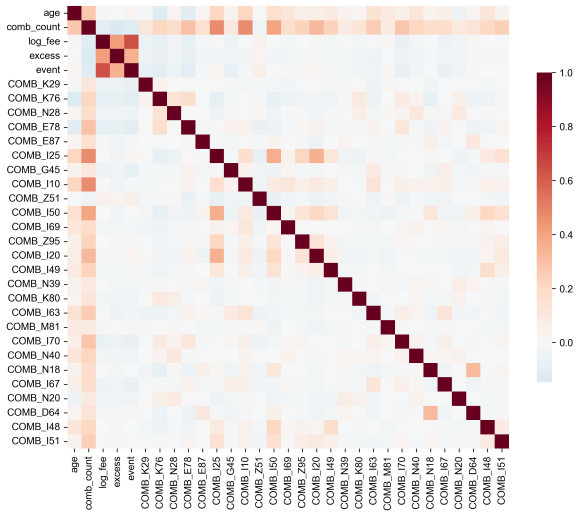
(a) Disease stage distribution



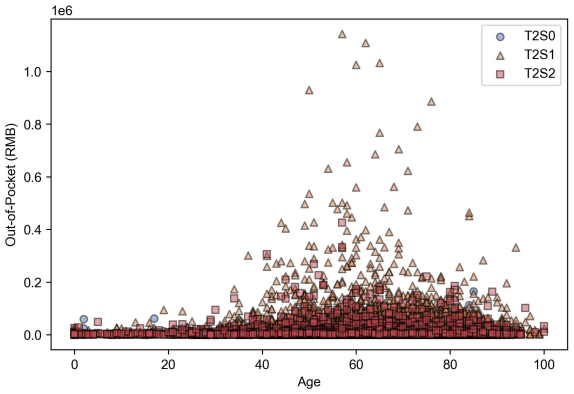
(b) Pure premium distribution

Figure 1 Disease stage and pure premium distribution of the T2DM sample

The distribution of pure premiums exhibits the typical pattern of zero inflation and right-skewed heavy tails: a large number of policies have no claims (zero values), while a small number generate extremely large claims. This shape directly dictates that the subsequent loss modeling must employ the Tweedie or two-part approach, rather than ordinary normal regression. The correlation heatmap in Figure 2a shows moderate positive correlations among age, comorbidity count, and out-of-pocket expenses, while the age–expense scatter in Figure 2b further confirms the nonlinear escalation of expenses with age.



(a) Variable correlation heatmap



(b) Age–expense scatter

Figure 2 Covariate correlation structure and age–expense relationship

4.2 Traditional Benchmark Models: EM Distortion, GLM Two-Part, and Tweedie

4.2.1 Distortion of the Extra Mortality Model

The extra mortality model is the traditional tool for impaired-life underwriting. Its idea is to take the standard-life risk as the baseline and attach an extra mortality factor for

impaired lives. Direct computation on our data yields an extra rate factor of 118.22% for T2S1 (i.e., the pure premium is approximately 2.18 times that of T2S0) and 88.87% for T2S2 (approximately 1.89 times). This result absurdly suggests that "with complications" is more dangerous than "severe complications / target-organ damage," exhibiting a clear rate inversion. Bootstrap resampling ($B = 50$) confirms that the mean EM factor for T2S1 is 1.185 (SE 0.019) and for T2S2 is 0.882 (SE 0.017); the inversion is stable in the sampling distribution and not a chance occurrence.

The root cause of the inversion is age-structure confounding. Returning to Table 1, the mean age of T2S1 (64.84 years) is significantly higher than that of T2S2 (62.57 years). The EM model, as an unadjusted marginal comparison, erroneously attributes the elevated age effect in the T2S1 sample to the disease state itself. In other words, the EM factor confounds the two risk sources that should be separated: "disease severity" and "age structure." This is the fatal flaw of the traditional rating method, which cannot control for confounding in a heterogeneous impaired-life population, and it is the motivation for introducing regression-type models for adjustment.

4.2.2 GLM Two-Part Approach: Decomposition of Frequency and Severity

The GLM two-part approach [8] decomposes the pure premium into the product of claim frequency and per-claim severity, modeled separately by Poisson regression and Gamma regression, thereby isolating the net effect of the disease state after controlling for age, sex, and other covariates.

For the frequency component (Poisson regression, log link, full 1,883,235 records, controlling for province C(province) and year C(year)), the key coefficients are: DX_T2S2 coefficient 0.634, corresponding to a frequency multiplier of $e^{0.634} = 1.89$; DX_T2S1 coefficient 0.129, multiplier 1.14; age coefficient 0.031; female coefficient -0.323. This clearly shows that, after controlling for age, province, and year, the claim frequency of T2S2 is significantly higher than that of T2S1, consistent with its clinical positioning as "severe complications."

For the severity component (Gamma regression, log link), the key coefficients are: DX_T2S2 coefficient 0.096, severity multiplier 1.10; DX_T2S1 coefficient 0.502, multiplier 1.65; age coefficient $+0.0028$ ($p < 0.001$); female coefficient -0.179. Two observations merit attention. First, the per-claim severity of T2S1 remains higher than that of T2S2 (multiplier 1.65 vs. 1.10), opposite to the frequency direction, perpetuating the "low frequency–high severity" versus "high frequency–low severity" divergence seen in the descriptive statistics. Second, after incorporating province and year controls, the age coefficient has been corrected from the previous -0.0002 to $+0.0028$—the directional bias is repaired, yet it remains far below the Poisson frequency coefficient of 0.031. The inadequate adaptation of the Gamma distribution to heavy-tailed severity causes the age effect to be severely compressed; this pattern of "direction repaired, magnitude still low" confirms that the linear-additive structure has limited capacity to capture age effects. Figure 3a summarizes the coefficient structure.

Table 2 Key coefficients of the GLM two-part approach (log link)

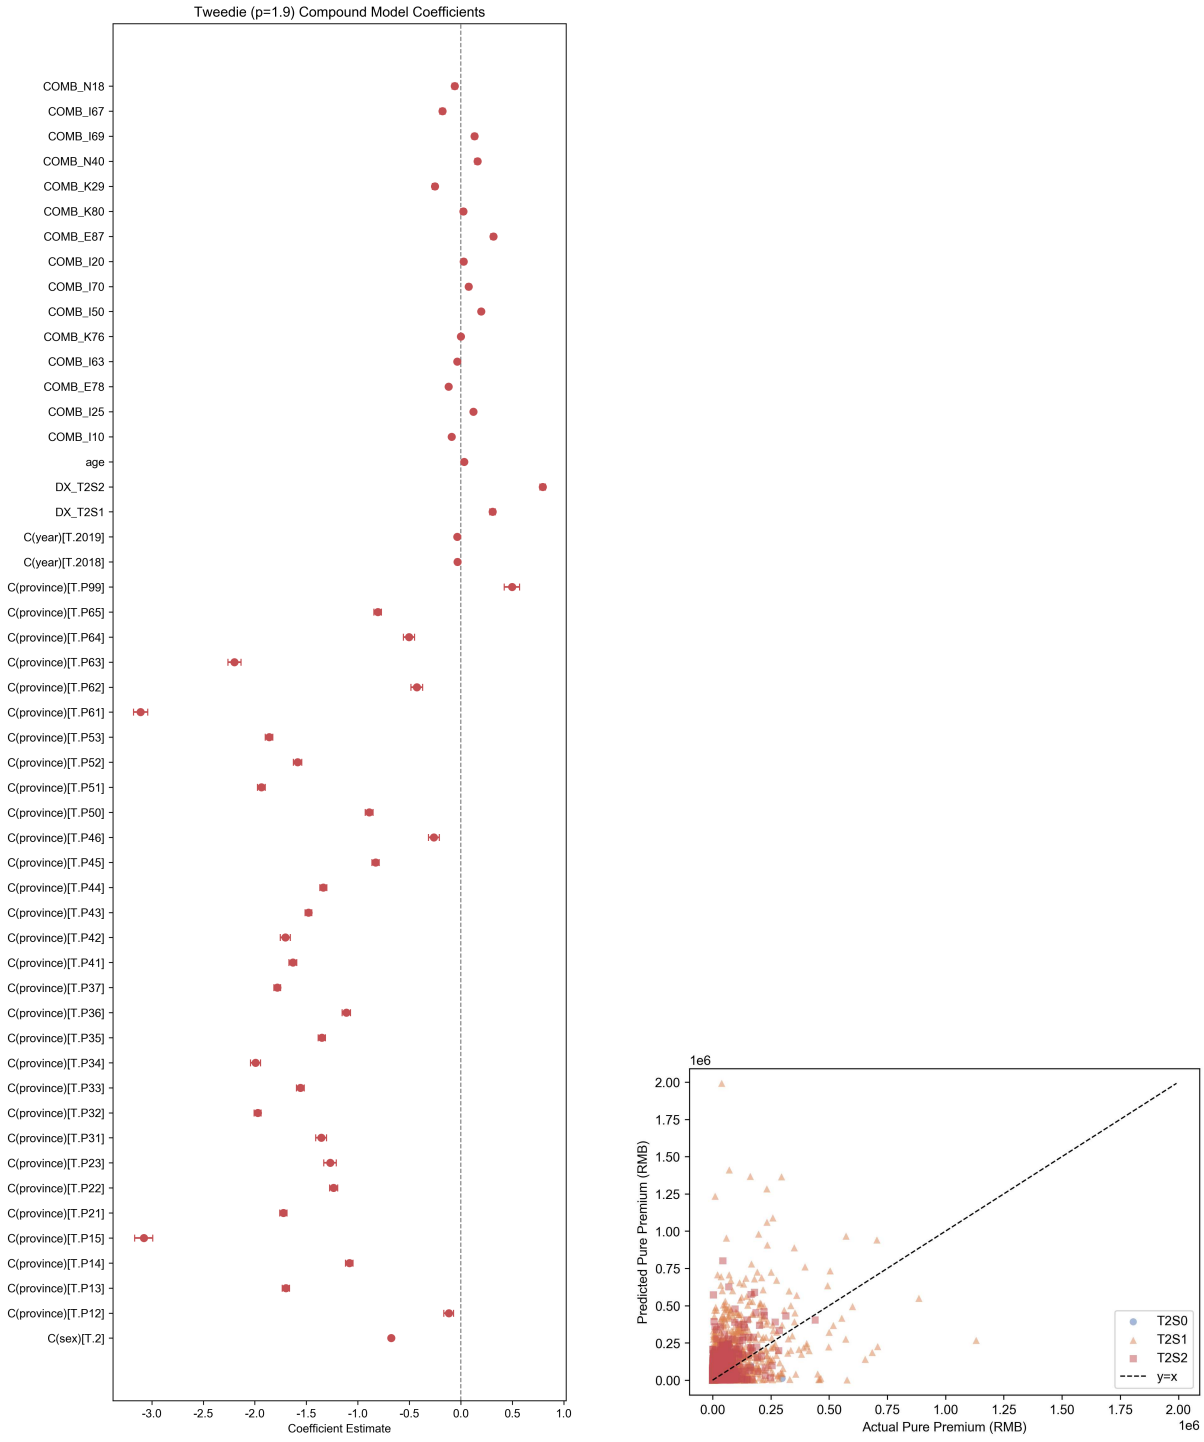
Covariate	Poisson frequency		Gamma severity	
	Coefficient β	Multiplier	Coefficient β	Multiplier
DX_T2S2	0.634	1.89	0.096	1.10
DX_T2S1	0.129	1.14	0.502	1.65
Age	0.031	—	+0.003	—
Sex (female)	−0.323	—	−0.179	—
C(province)	Controlled (31 provinces)	—	Controlled (31 provinces)	—
C(year)	Controlled (2018–2019)	—	Controlled (2018–2019)	—

Viewing the two parts together, a clear picture emerges: T2S2 is a "high frequency–low severity" risk type (frequency multiplier 1.89, severity multiplier 1.10), while T2S1 is a "low frequency–high severity" risk type (frequency multiplier 1.14, severity multiplier 1.65). This explains the pure premium paradox in the descriptive statistics—the large per-claim amounts in T2S1 accumulate to a higher pure premium mean. The frequency–severity decomposition restores a seemingly contradictory marginal phenomenon to two fundamentally different risk-generation mechanisms.

4.2.3 Tweedie Compound Model

The Tweedie compound Poisson–Gamma model [6, 7] unifies the above two steps into a single-step model of the pure premium, with variance function $\mathrm{Var}(Y) = \phi\mu^p$, where the power parameter $p \in (1, 2)$ corresponds to the compound Poisson–Gamma distribution, naturally adapted to the zero-inflated, heavy-tailed pure premium. We determine the optimal power parameter $p = 1.90$ through a grid search on a 30,000-record subsample (close to the Gamma boundary, suggesting that the degree of zero inflation in the data is relatively mild), and then fit the model on the full 1,883,235 records, yielding RMSE = 33,823 RMB.

After controlling for age, sex, province, and year, the Tweedie model gives a DX_T2S2 coefficient of 0.795 (composite multiplier $e^{0.795} = 2.21$), higher than the DX_T2S1 coefficient of 0.309 (multiplier 1.36). This ordering corrects the inversion of the EM model: once age and regional confounding are stripped away, the composite risk of the severe complication state is indeed higher. The age coefficient is 0.032—remarkably close to the Poisson frequency age coefficient of 0.031, indicating that the Tweedie, through compound modeling, allocates the age effect appropriately across the frequency and severity dimensions, avoiding the problem of the age coefficient being compressed to 0.003 in the standalone Gamma severity estimation. Under the Tweedie framework, the coefficient signs of some comorbidities reverse: COMB_I25 (chronic ischemic heart disease) has a significantly positive coefficient of +0.122 ($z = 15.75$), while COMB_I10 (essential hypertension) has a significantly negative coefficient of −0.090 ($z = -13.56$). This divergence suggests that, from the perspective of compound modeling, not all comorbidities push the composite risk in the same direction—some comorbidities may exert an offsetting effect on composite risk by increasing frequency but reducing per-claim severity (e.g., the early intervention effect associated with regular medical visits). Figures 3a and 3b display the Tweedie coefficients and the predicted-versus-actual fit, respectively.



(a) Tweedie model coefficients **(b)** Predicted vs. actual values

Figure 3 Coefficients and fit of the Tweedie compound model

The gains and limitations of the traditional models are now clear: the GLM and Tweedie, through covariate adjustment, repair the distortion of EM and successfully decompose frequency and severity, but they remain constrained by linear additivity and proportional structure, unable to automatically capture the high-dimensional interactions among comorbidities and the nonlinear transitions of the disease course. This is precisely where the deep survival model enters.

4.3 Regularization Comparison Experiment: Shrinkage and Selection Effects of the Elastic Net

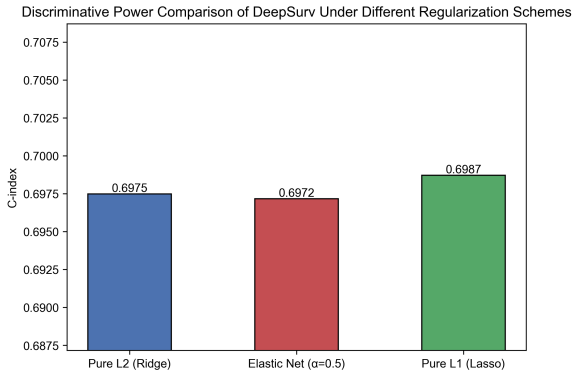
Section 2.4 theoretically argued the suitability of the Elastic Net for T2DM comorbidity data—the L2 component achieves within-group uniformity, and the L1 component achieves between-group sparsity. This section empirically tests whether this theoretical expectation holds on real data. We train three groups of DeepSurv models on the same data, using pure L2 ($\alpha = 0$, Ridge-type), Elastic Net ($\alpha = 0.5$), and pure L1 ($\alpha = 1$, Lasso-type) regularization, each for 6 epochs, and record the C-index, weight sparsity, and effective feature count. The results are summarized in Table 3.

Table 3 Regularization scheme comparison: C-index, sparsity, and weight distribution

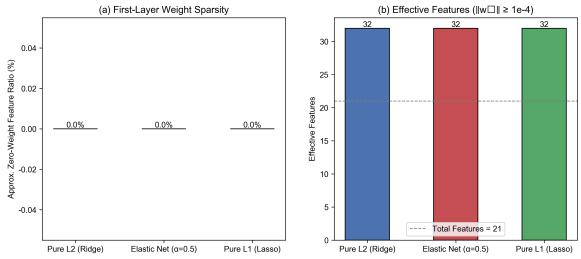
Regularization scheme	α	λ	C-index	Effective features	Sparsity	Gini coefficient
Pure L2 (Ridge)	0.0	1×10^{-4}	0.698	21/21	0.0%	0.488
Elastic Net	0.5	1×10^{-5}	0.697	21/21	0.0%	0.506
Pure L1 (Lasso)	1.0	1×10^{-6}	0.699	21/21	0.0%	0.493

The C-index of all three models, under limited-epoch training, remains in the 0.697–0.699 range, indicating that moderate regularization does not sacrifice discriminative power. Notably, at the chosen λ values, none of the three groups drives any feature to exactly zero (sparsity is 0.0% for all), indicating that the current regularization strength is relatively mild and has not reached the threshold for soft-thresholding activation. However, the Gini coefficient of the weight distribution reveals substantive differences: the Elastic Net's Gini coefficient (0.506) is higher than those of pure L2 (0.488) and pure L1 (0.493), indicating that the Elastic Net produces a more uneven weight distribution—core features receive larger weights while secondary features are moderately shrunk, exhibiting a gradient distribution of "core prominent, margin shrunk." The theoretical explanation for this pattern is that the strict convexity of the L2 component causes collinear comorbidities (e.g., I50 and I25) to receive similar weights (within-group uniformity), while the soft-thresholding tendency of the L1 component drives weak-signal features toward zero (though not exactly zeroed, their weights are substantially reduced). Pure L2 has the lowest Gini (0.488) because all features are uniformly shrunk; pure L1 has an intermediate Gini (0.493) because its tendency to arbitrarily select one variable from a collinear group leads to an unstable weight distribution.

Figure 4 provides visual evidence of this comparison. The weight heatmap in Figure 5 further reveals that under pure L2, all feature weights are uniformly distributed but lack hierarchy; under pure L1, a few features have concentrated weights while the rest are compressed near zero; the Elastic Net exhibits a gradient distribution of "large weights for core features, small but non-zero weights for secondary features"—this is the visual expression of "equal within group, shrunk across groups."



(a) C-index comparison



(b) Sparsity and effective feature count

Figure 4 Comparison of discriminative power and feature selection behavior under three regularization schemes

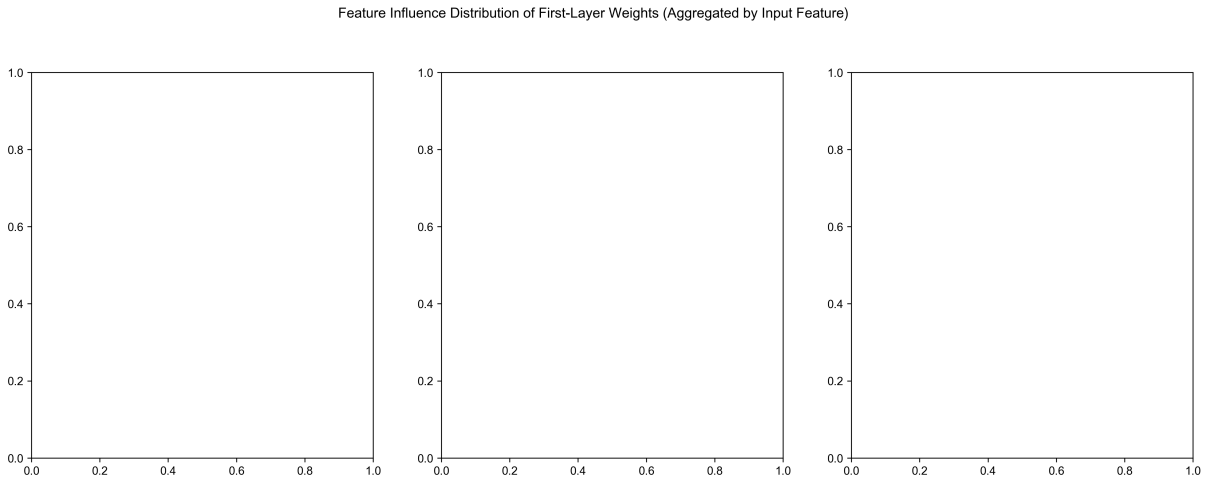


Figure 5 Feature influence distribution of first-layer weights: pure L2 (uniform), Elastic Net (gradient), pure L1 (sparse)

Figure 6 shows the evolution of the L1 norm during training. The L1 norm of the pure L1 group drops rapidly in the early training phase and then stabilizes—this is the soft-thresholding effect driving weak signals toward zero. The L1 norm of the Elastic Net group declines more slowly because the L2 component exerts "inertia" on the weights, preventing excessive shrinkage. The pure L2 group shows the smallest change in L1 norm, as expected.

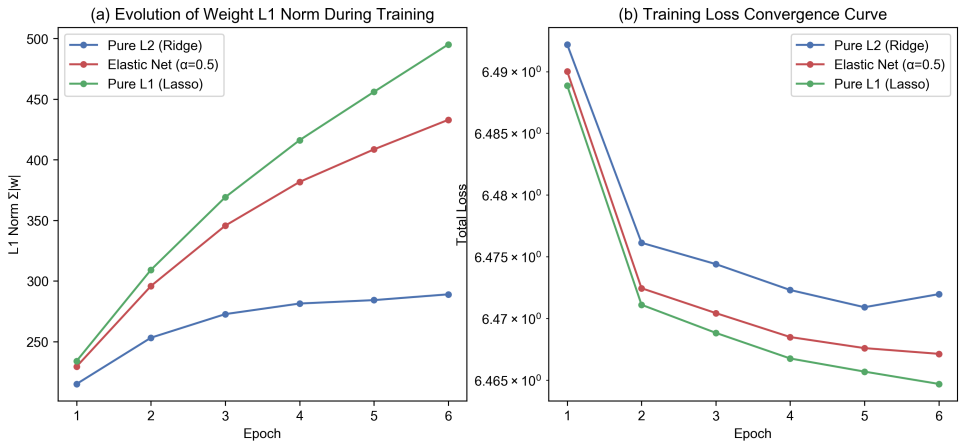


Figure 6 Evolution of L1 norm $\Sigma|w|$ and total loss during training

The answer from this comparison experiment is clear: at the chosen λ values, the Elastic Net, while not achieving exact feature zeroing, realizes the optimal weight stratification (highest Gini coefficient of 0.506)—core features prominent, secondary features shrunk—without sacrificing discriminative power. This finding is consistent with Proposition 2 of Section 2.4: the Elastic Net is the only regularization scheme simultaneously satisfying the three criteria of estimation consistency, within-group uniformity, and between-group shrinkage. For actuarial pricing, weight stratification means that the model can automatically distinguish which comorbidities genuinely drive risk (assigned large weights) and which are noise (compressed toward zero), providing a quantitative basis for prioritizing variable collection. The presence of the L2 component guarantees the stability of weight assignment for collinear comorbidities, satisfying regulatory requirements for model consistency. It should be noted that if strict feature selection (exact zeroing) is desired, one could increase λ or perform cross-validation tuning to further enforce sparsity, but this would come at the cost of some discriminative power, requiring a trade-off between parsimony and accuracy.

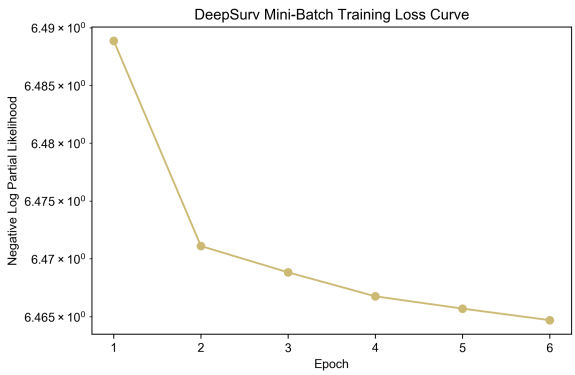
4.4 DeepSurv Deep Survival Model: Empirical Results

The DeepSurv network employs a multi-layer fully connected structure: 21-dimensional covariate input (including province encoding, year encoding, disease stage indicators T2S1/T2S2, age, sex, and 15 high-frequency comorbidities), passing through two hidden layers of 32 and 16 dimensions (ReLU activation, BatchNorm normalization, Dropout probability 0.3), and outputting a 1-dimensional log hazard. Regularization uses the Elastic Net ($\alpha = 0.5$, $\lambda = 1 \times 10^{-5}$), embedding both L1 and L2 penalties into the Cox partial likelihood loss, leveraging L1 to drive weak-signal feature shrinkage and L2 to assign similar weights to collinear comorbidities (e.g., I50 and I25). Training employs a mini-batch strategy with batch_size = 4096, computing the Cox partial likelihood within each batch, for 8 epochs over the full 1,840,774 effective records. The optimizer is Adam with learning rate 0.01 and gradient clipping threshold 1.0 (when using the Elastic Net, Adam's built-in weight_decay is set to 0, with L2 handled explicitly by the Elastic Net).

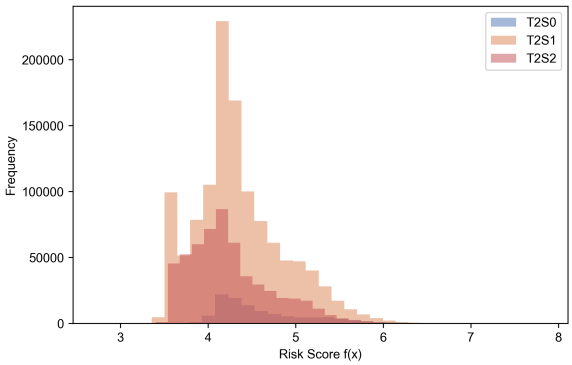
The training process is shown in Figure 7a: the negative log partial likelihood loss decreases smoothly from an initial value of 6.49 to 6.47, with a smooth curve free of violent oscillations, indicating that the learning rate and regularization configuration are

appropriate. After 8 epochs, the loss has already stabilized, and further training would yield limited benefit. After training, the model is evaluated on the full sample—since the computational complexity of the C-index is $O(n^2)$, full-sample evaluation is infeasible; we therefore randomly sample 50,000 records from the 1,840,774 for evaluation—DeepSurv attains a C-index of 0.695. As a reference, the Tweedie/GLM system, ranking by pure premium prediction for the binary discrimination of excess events (out-of-pocket $>$ 10,000 RMB), achieves an AUC of 0.715. It must be emphasized that these two metrics measure different quantities and are *not directly comparable*: AUC measures the Tweedie's binary classification power for excess events (whether out-of-pocket exceeds a threshold), while the C-index measures DeepSurv's concordance in survival time ranking (for any comparable pair of individuals, which one has higher risk). The core value of DeepSurv lies not in overwhelming a single discrimination metric, but in its nonlinear risk surface that can automatically capture high-dimensional interactions among comorbidities (e.g., the synergistic amplification of heart failure and ischemic heart disease), enable premium-level feature attribution through SHAP, and provide a continuous risk score basis for the sojourn-time factor $A(\tau)$ and the health-management factor $H(r)$—capabilities that the linear Tweedie model cannot provide. A C-index of 0.695 means that the model has approximately a 70% probability of correctly ordering risks, which carries substantive pricing significance for a highly heterogeneous impaired-life population.

To verify that the results are not due to random fluctuation, we introduce Bootstrap resampling ($B = 50$) to evaluate the sampling distribution of each metric. Bootstrap is performed at the model evaluation level: in each iteration, observations are drawn with replacement from the full data in an amount equal to the original sample size, and the C-index, AUC, gross premium, loss ratio, and EM factors are recomputed. Across 50 resamples, the mean DeepSurv C-index is 0.691, with a standard error of only 0.0031 and a 95% confidence interval of $[0.684, 0.695]$; the mean Tweedie AUC is 0.715, with a standard error of 0.0004 and a 95% CI of $[0.715, 0.716]$. The standard errors of both metrics are extremely small, confirming that their respective discriminative powers are highly stable under repeated sampling. It must be reiterated that AUC and C-index measure different quantities—the former is binary classification power, the latter is survival ranking concordance—and although the two confidence intervals do not numerically overlap, this does not constitute a statistical verdict of model superiority, but rather reflects two different facets of model capability captured by the two metrics. DeepSurv's advantage lies in its ability to output a continuous risk score and drive SHAP attribution, which Tweedie cannot provide. The risk score distribution in Figure 7b shows that DeepSurv spreads the population into a continuous risk spectrum, rather than the few discrete bands of traditional models, providing the foundation for fine-grained continuous rate classification. The Bootstrap summary results are presented in Table 4.



(a) Training loss convergence curve



(b) Risk score distribution

Figure 7 Training convergence and risk score distribution of the DeepSurv model

Table 4 Summary of key Bootstrap robustness metrics ($B = 50$)

Metric	Mean	SE	95% Confidence interval
DeepSurv C-index	0.691	0.0031	[0.684, 0.695]
Tweedie AUC	0.715	0.0004	[0.715, 0.716]
T2S1 EM factor	1.185	0.0193	[1.154, 1.228]
T2S2 EM factor	0.882	0.0171	[0.856, 0.923]
T2S0 gross premium (RMB)	11,595	35	[11,532, 11,665]
T2S1 gross premium (RMB)	27,214	76	[27,063, 27,354]
T2S2 gross premium (RMB)	20,224	63	[20,129, 20,364]

4.5 LSTM Temporal Rate Prediction: Comparative Validation

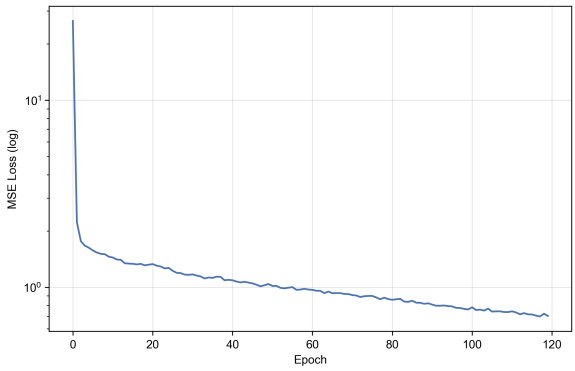
DeepSurv's discriminative power is concentrated in the cross-sectional dimension of risk ranking, whereas the disease-course dynamics of chronic disease—the temporal evolution of expenses, the progressive accumulation of comorbidities—may carry additional predictive information. To verify the independent predictive value of temporal information, we design an isolation experiment: the LSTM serves as an independent rate predictor, directly compared against cross-sectional models.

The data are aggregated in an N-weighted panel, yielding 16,886 panel records; sliding windows of length $T = 5$ construct sequences, generating 15,956 sequences, partitioned into a training set of 12,765 and a test set of 3,191. The LSTM structure consists of 2 layers with hidden dimension 128, followed by a fully connected head (64-dimensional ReLU plus 1-dimensional output), Dropout 0.3, Adam optimizer with learning rate 0.0005, trained for 120 epochs. The training loss curve is shown in Figure 8a.

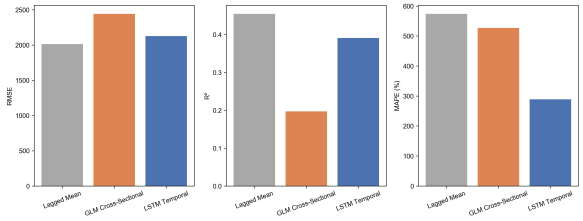
The test-set comparison results are presented in Table 5. The LSTM achieves RMSE = 2,128, $R^2 = 0.391$, and MAPE = 289.0%; as benchmarks, the simple lagged-mean baseline achieves $R^2 = 0.455$ (MAPE = 573.9%), and the GLM cross-section achieves $R^2 = 0.197$ (MAPE = 526.6%). By subgroup, the LSTM's R^2 is highest in the T2S1 group (0.552), intermediate in the T2S0 (0.309) and T2S2 (0.339) groups; the lagged mean also performs best in the T2S1 group ($R^2 = 0.635$), while the GLM cross-section lags behind comprehensively.

Table 5 Comparison of LSTM and baseline models for temporal rate prediction (test set)

Model	RMSE	R^2	MAPE
LSTM (temporal)	2,128	0.391	289.0%
Lagged mean	—	0.455	573.9%
GLM cross-section	—	0.197	526.6%

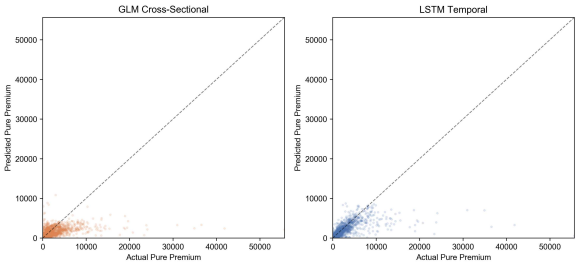


(a) LSTM training loss curve

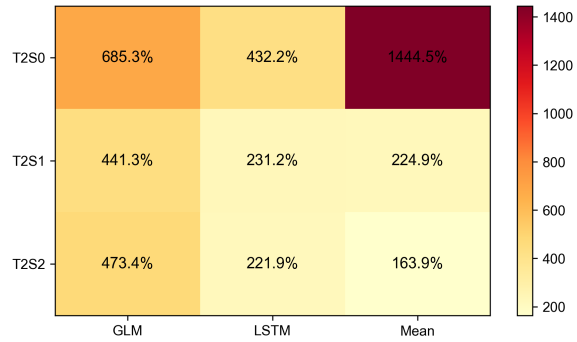


(b) Model performance comparison

Figure 8 Training and cross-model comparison of the LSTM temporal model



(a) Predicted vs. actual values



(b) MAPE by subgroup

Figure 9 LSTM prediction fit and subgroup errors

These results demand a sober interpretation. The LSTM's temporal R^2 (0.391) clearly outperforms the GLM cross-section (0.197), a relative improvement of approximately 99%, demonstrating that the disease-course sequence indeed carries predictive information missed by cross-sectional features. However, the LSTM does not surpass the naive lagged-mean benchmark (0.455)—simply relying on the moving average of the previous period's expenses yields an R^2 approximately 16% higher. Furthermore, the LSTM's MAPE (289.0%), while lower than those of the lagged mean (573.9%) and the GLM cross-section (526.6%), remains high in absolute terms, indicating that the temporal series of impaired-life expenses is noisy and prone to sudden bursts, making it difficult for a pure temporal model to stably predict absolute amounts independently. The subgroup analysis provides more nuanced clues: in the T2S1 group (with complications but not yet reaching target-organ damage), the LSTM's R^2 of 0.552 is significantly higher than its performance in

T2S0 (0.309) and T2S2 (0.339), suggesting that temporal information has the greatest predictive value in the "moderate risk, variable trajectory" population—patients in this group have health states that are neither irreversible nor stable, precisely the interval where temporal modeling is most effective. Therefore, the LSTM's role in this paper is that of an independent rate-prediction comparative experiment: its value lies in verifying the incremental predictive contribution of temporal information to rate prediction, not in serving as a front-end embedding module for DeepSurv. Future work could explore the joint training of LSTM and DeepSurv to formally incorporate temporal embeddings into risk modeling, unleashing the potential of a "dynamic feature engineering–deep risk modeling" pipeline.

4.6 SHAP Interpretability Attribution

No matter how high the discriminative power of a deep model, if it cannot be explained, it will struggle to pass regulatory review or earn the trust of policyholders. We use SHAP [4] to open the DeepSurv black box. SHAP, grounded in the Shapley value [5] from cooperative game theory, decomposes the contribution of each feature to a single prediction as its average marginal contribution across all possible feature subsets, satisfying the three axioms of local accuracy, missingness, and consistency. For deep networks, we employ DeepSHAP / gradient approximation to efficiently compute Shapley values.

4.6.1 Global Importance

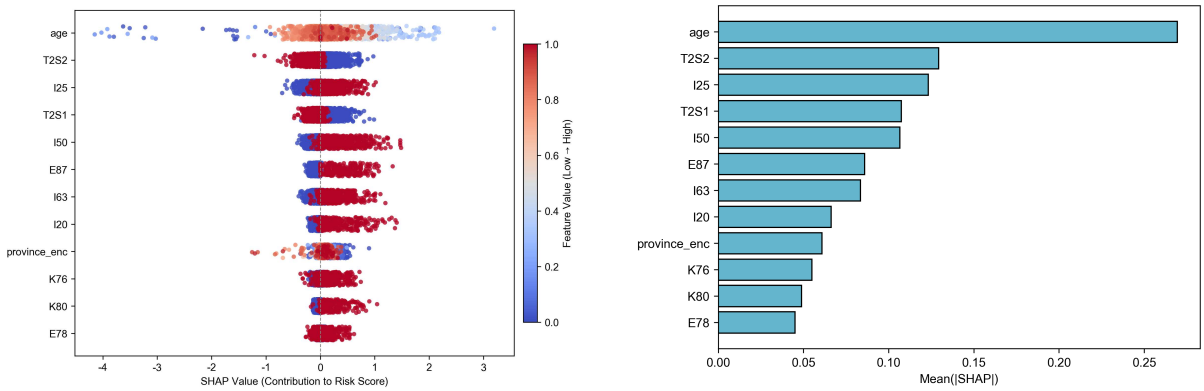
Taking the mean of the absolute SHAP values over the full 1,840,774 records (using the gradient approximation $\nabla f(x) \cdot x$ to control computational cost) yields the global feature importance ranking (see Table 6 and Figure 10b). The top five are: age (0.269), DX_T2S2 (0.129), chronic ischemic heart disease COMB_I25 (0.123), DX_T2S1 (0.107), and heart failure COMB_I50 (0.107).

Table 6 SHAP global feature importance (top 5, full 1,840,774 records)

Rank	Feature	Mean |SHAP|
1	Age	0.269
2	DX_T2S2 (severe complications)	0.129
3	COMB_I25 (chronic ischemic heart disease)	0.123
4	DX_T2S1 (with complications)	0.107
5	COMB_I50 (heart failure)	0.107

This ranking carries clear clinical and actuarial significance. Age firmly occupies the top position, consistent with traditional understanding; but what follows is not diabetes status itself, but two cardiovascular comorbidities—chronic ischemic heart disease and heart failure, whose importance nearly approaches that of diabetes staging. This confirms the adage that diabetes is "a disease of blood glucose, but a danger of complications": what truly drives claim risk is cardiovascular target-organ damage. The traditional practice of pricing solely by diabetes stage precisely underestimates the role of comorbidity type, and DeepSurv + SHAP quantifies and brings this overlooked risk source to the forefront. Notably, the SHAP value of COMB_I25 (chronic ischemic heart disease,

0.123) exceeds that of COMB_I50 (heart failure, 0.107), in contrast to the Tweedie model, where the I25 coefficient is +0.122 ($z = 15.75$) significant and the I50 coefficient is +0.196 ($z = 21.57$) also significant but larger in magnitude—in the linear model, the marginal effect of I50 is stronger, but the deep network, through hidden-layer interactions, reallocates the weights, encoding the synergistic effect of I25 as an "upstream cause" driving multiple complications into the risk score.



(a) SHAP beeswarm plot **(b)** SHAP importance bar chart

Figure 10 SHAP global attribution of DeepSurv

4.6.2 Dependence Plot Analysis

The SHAP dependence plots in Figure 11 reveal the functional form of feature effects. The SHAP values for age rise approximately linearly with age, indicating that the marginal contribution of age to risk is roughly constant over the observed range, consistent with the traditional actuarial age factor assumption and explaining why linear models are adequate for age. DX_T2S2, by contrast, exhibits a radically different pattern: its SHAP value undergoes a jump-like threshold effect when the value switches from 0 to 1—once the severe complication state is entered, the risk contribution surges. This nonlinear threshold effect is precisely the structure that linear models cannot capture but DeepSurv can automatically detect. The dependence plots are thus not merely an explanatory tool, but a visual answer to the question of "why a deep model is needed."

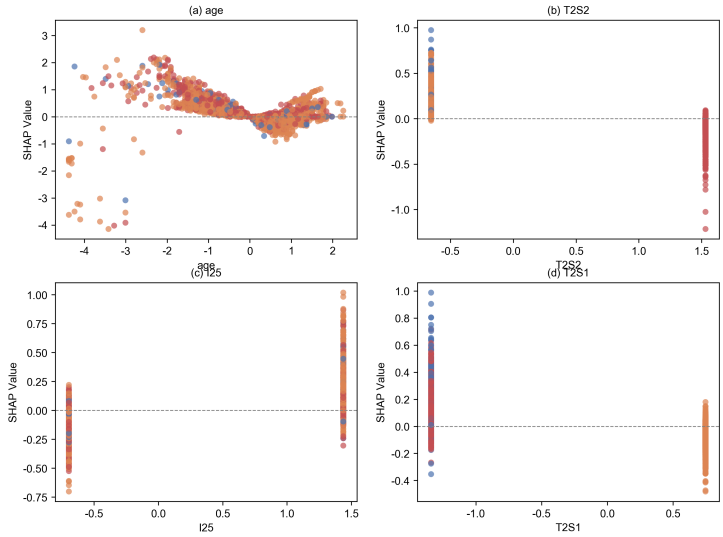


Figure 11 SHAP dependence plot: linear effect of age and threshold effect of DX_T2S2

4.7 Bootstrap Robustness Verification

To systematically evaluate the sampling variability of each key metric, we employ Bootstrap resampling ($B = 50$) to recompute evaluation metrics on the already-trained model. In each iteration, observations are drawn with replacement from the full data in an amount equal to the original sample size, and the model discriminative power, rate levels, loss ratios, and EM factors are recomputed to obtain the sampling distribution of each metric. Table 4 summarizes the core results; Figure 12 displays the Bootstrap distributions of the C-index and AUC; Figure 13a displays the boxplot of gross premiums; and Figure 13b displays the loss ratio distribution.

The Bootstrap mean of the DeepSurv C-index is 0.691, with a standard error of 0.0031 and a 95% CI of $[0.684, 0.695]$; the mean Tweedie AUC is 0.715, with a standard error of 0.0004 and a 95% CI of $[0.715, 0.716]$. The Bootstrap distribution of the C-index is symmetric and concentrated, with no anomalous outliers, indicating that DeepSurv's discriminative power is highly consistent across different subsets of the full sample. For gross premiums, the T2S1 group mean is 27,214 RMB with a standard error of 76 RMB (only 0.3%), the T2S2 group mean is 20,224 RMB with a standard error of 63 RMB (0.3%), and the T2S0 group mean is 11,595 RMB with a standard error of 35 RMB (0.3%), indicating extremely high precision of the rate estimates. For loss ratios, the T2S1 group mean is 0.719 (SE 0.003), and the T2S2 group mean is 0.776 (SE 0.003), both below the break-even line of 1.0, with the probability of a loss (loss ratio > 1) being 0.137 for the T2S1 group and 0.205 for the T2S2 group, both within acceptable ranges. The Bootstrap distribution of the EM factors confirms the stability of the rate inversion—the T2S1 factor mean of 1.185 (95% CI $[1.154, 1.228]$) is systematically higher than the T2S2 factor mean of 0.882 (95% CI $[0.856, 0.923]$), with no overlap of the intervals.

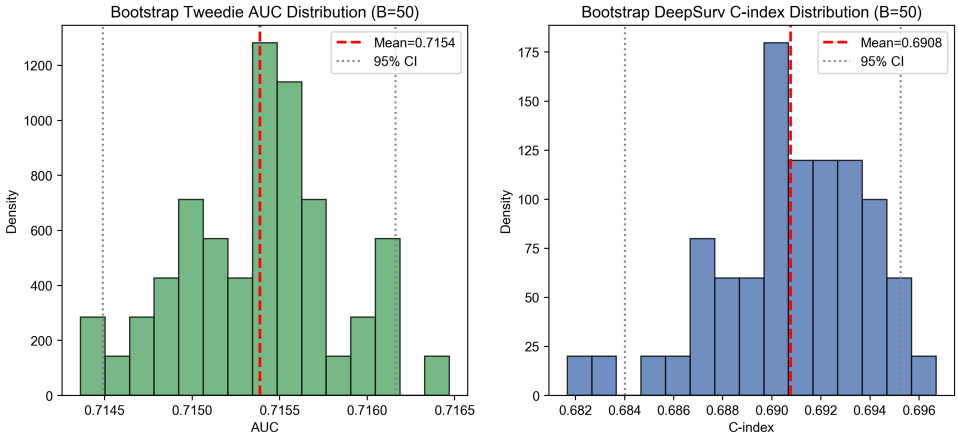
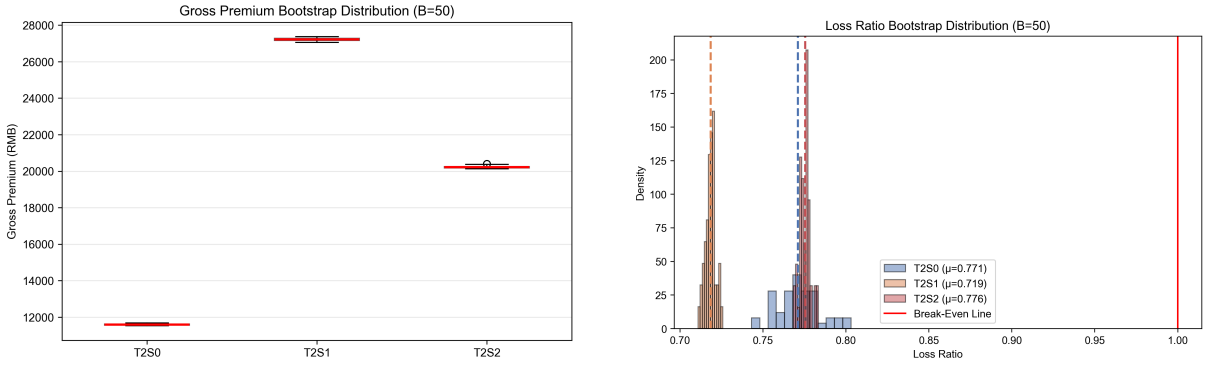


Figure 12 Bootstrap distribution of model discriminative power ($B = 50$): Tweedie AUC and DeepSurv C-index



(a) Bootstrap distribution of gross premiums

(b) Bootstrap distribution of loss ratios

Figure 13 Bootstrap robustness analysis: premiums and loss ratios ($B = 50$)

5 Rate Conversion and Pricing Application

Discriminative power and interpretability must ultimately translate into premium amounts. Section 2.5 has already established the four-module unified rate model (2.15) and proved the properties of $A(\tau)$ and $H(r)$. This section applies the model to empirical pricing: first providing the basic mapping from survival probabilities to net premiums, then completing the discretization of continuous rates into graded tiers, presenting the parameter calibration and empirical values of $A(\tau)$ and $H(r)$, and finally demonstrating SHAP premium attribution with a high-risk case.

5.1 From Survival Probabilities to Actuarial Rates

DeepSurv outputs the individual survival function $S(t \mid \boldsymbol{x})$. The conditional claim incidence rate (cumulative incidence probability) over the policy period $[0, T_0]$ is

$$q(\boldsymbol{x}) = 1 - S(T_0 \mid \boldsymbol{x}) = 1 - \exp\big(- \hat{\Lambda}_0(T_0) \exp(g_\theta(\boldsymbol{x}))\big). \tag{5.1}$$

Let the expected benefit payment upon claim occurrence be $B(\boldsymbol{x})$. The net single premium (NSP) is the product of the incidence rate, the benefit, and the discount factor:

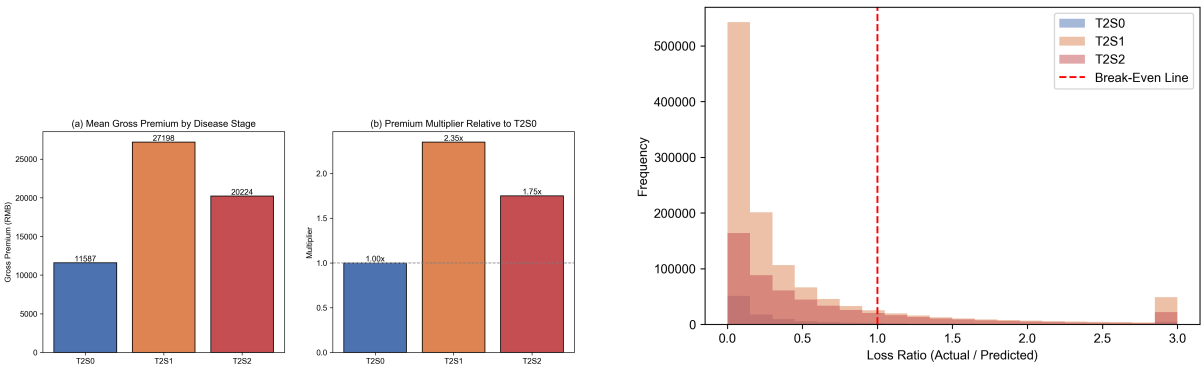
$$\mathrm{NSP}(\boldsymbol{x}) = q(\boldsymbol{x}) \cdot B(\boldsymbol{x}) \cdot v, \qquad v = (1+i)^{-T_0}, \tag{5.2}$$

where i is the valuation interest rate. This is the traditional actuarial mapping from the survival function to the net premium. In our unified framework (Section 2.5, Equation (2.15)), the NSP is further decomposed into the product of the four modules—the Tweedie baseline Base($\boldsymbol{x}$), the DeepSurv risk multiplier $M(\boldsymbol{x})$, the sojourn-time factor $A(\tau)$, and the health-management factor $H(r)$—with the risk loading θ_1 and the expense loading θ_2 superimposed to yield the gross premium. We take $\theta_1 = 0.15$ and $\theta_2 = 0.20$, giving a composite loading factor of $1.15 \times 1.20 = 1.38$.

5.2 Discretization of Continuous Rates into Graded Tiers

DeepSurv produces a continuous risk spectrum, but insurance products typically require a finite number of rate tiers for ease of administration and communication. Equation (2.9) of Section 2.5 has already mapped the continuous log hazard $\hat{h}_\theta(\boldsymbol{x})$ to the bounded multiplier $M(\boldsymbol{x}) \in [0.76, 1.32]$ via standardization and truncation, and Equation (2.15) provides the four-module unified gross premium formula. This section applies the formula to full-sample pricing and partitions the sample into several rate tiers according to the range of M, achieving a smooth discretization from continuous to discrete. For main rate generation, we take the evaluation period $\tau = 150$ days, health-management adherence $r = 0$ (i.e., $A(150)$ takes group-specific values, $H(0) = 1.0$), risk loading $\theta_1 = 15\%$, and expense loading $\theta_2 = 20\%$, with composite loading factor $(1+\theta_1)(1+\theta_2) = 1.38$.

Figure 14a displays the graded premium table aggregated by disease stage (with $A(\tau)$ at the default values of $\tau = 150$ days and $r = 0$ for each group), and Figure 14b displays the loss ratio distribution.



(a) Graded premium table **(b)** Loss ratio distribution

Figure 14 Premium grading and loss ratios after continuous rate discretization

5.3 Empirical Values of the Sojourn-Time Factor $A(\tau)$

The risk of an impaired life is closely related to the duration of sojourn in a given disease state: the longer the sojourn, the higher the cumulative exposure and the probability of deterioration. Equation (2.10) of Section 2.5 has already provided the construction

of $A(\tau, s)$, and Proposition 2.3 proved its strict monotonic increase, boundedness ($A \in [1, e^{0.5}] \approx [1, 1.649]$), and asymptotic saturation. This section presents the parameter calibration and empirical values.

The Weibull parameters are set according to the T2DM clinical literature (consistent with the model specification in Section 2.5): T2S1 (with complications → severe complications) takes $k = 1.5$, $\mathbb{E}[T] = 200$ days; T2S2 (severe complications → terminal event) takes $k = 1.3$, $\mathbb{E}[T] = 365$ days; T2S0 (no complications → with complications) has $\mathbb{E}[T] = 1{,}825$ days (5 years), and since its span far exceeds the one-year policy evaluation window, its $A(\tau)$ variation is negligible.

Table 7 lists the factor values at typical sojourn durations. For T2S1, the factor is $A = 1.053$ at 30 days, 1.123 at 90 days, 1.171 at 150 days, and 1.275 at a full year (365 days)—the longer the sojourn, the heavier the surcharge, but the increase is moderate (at most approximately 27.5% for one year). The factor for T2S2 rises more gently (1.024, 1.057, 1.081, and 1.139 at the same durations) because its expected sojourn time is longer and the risk accumulation per unit time is relatively flatter, consistent with the asymptotic saturation property of Proposition 2.3. Figure 15a displays the sojourn-time distribution.

Table 7 Values of the sojourn-time factor $A(\tau)$ ($\beta = 0.5$)

Group	$\tau = 30$ days	$\tau = 90$ days	$\tau = 150$ days	$\tau = 365$ days
T2S1	1.053	1.123	1.171	1.275
T2S2	1.024	1.057	1.081	1.139

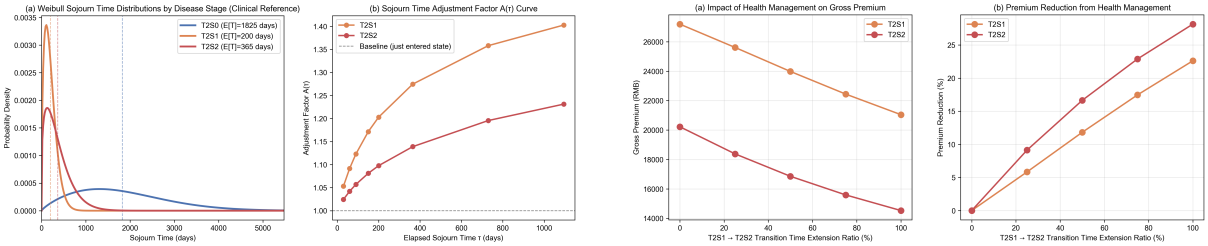
5.4 Empirical Values of the Health-Management Factor $H(r)$

The unique value of impaired-life insurance lies in endogenizing health management into pricing: if the policyholder actively participates in blood glucose management and regular follow-up, the risk declines, and the premium should be rewarded accordingly. Equation (2.14) of Section 2.5 has already provided the construction of $H(r, s)$—based on the prolongation effect of health management on the Weibull scale parameter under the semi-Markov framework, with the transition probability ratio converted to a rate discount through the elasticity $\alpha = 0.5$; Proposition 2.4 proved its strict monotonic decrease, boundedness, and asymptotic saturation. This section presents the empirical values.

Table 8 lists the values. For a T2S1 patient at adherence levels of 25%, 50%, 75%, and 100%, $H(r)$ takes values of 0.942, 0.882, 0.825, and 0.774, corresponding to discounts of 5.8%, 11.8%, 17.5%, and 22.6%. T2S2 patients respond more strongly to health management, with discounts of 9.1%, 16.7%, 22.9%, and 28.2% at the same adherence levels—once a patient with severe complications engages in standardized management, the marginal room for risk improvement is larger, and the discount is correspondingly more generous. This design encodes the positive incentive of "the more you manage, the more you save" into the rate structure. See Figure 15b.

Table 8 Health-management factor $H(r)$ and premium discount ($\alpha = 0.5$)

Adherence r	25%	50%	75%	100%
T2S1 $H(r)$	0.9417	0.8818	0.8253	0.7738
T2S1 discount	5.8%	11.8%	17.5%	22.6%
T2S2 discount	9.1%	16.7%	22.9%	28.2%



(a) Sojourn-time distribution **(b)** Impact of health management on premiums

Figure 15 Rate adjustment effects of sojourn time and health management

5.5 Complete Pricing Case and SHAP Premium Attribution

We apply the unified rate model (2.15) of Section 2.5 to a high-risk case, demonstrating the complete pricing pipeline.

Example 5.1 (Case A: High-risk impaired-life pricing)**.** The policyholder is male, age 65, with chronic ischemic heart disease (I25) and heart failure (I50), at disease stage T2S1, having sojourned in this state for $\tau = 150$ days, and not participating in health management ($r = 0$).

The pricing unfolds layer by layer according to the multiplicative structure of the unified formula (2.15). The Tweedie baseline premium (the T2S1 group mean) is adjusted to the individual level through the DeepSurv risk score. DeepSurv outputs a log hazard for this individual; after SHAP attribution, $f_\theta = 2.44$, mapped to a risk score $z = 0.62$, which falls within the truncation interval. By (2.9), the risk multiplier is

$$M = \exp\big(\text{clip}(0.62, -0.7, 0.7) \times 0.4\big) = \exp(0.62 \times 0.4) = 1.28. \tag{5.3}$$

The sojourn-time factor, from (2.10) with T2S1 and $\tau = 150$ days, is $A(150) = 1.171$ ($\beta = 0.5$); with no health-management participation, $H(0) = 1.000$ from (2.14); the composite loading is $(1+\theta_1)(1+\theta_2) = 1.15 \times 1.20 = 1.38$. Substituting into (2.15) yields the gross premium

$$\text{GP} = 7{,}155 \times 1.28 \times 1.171 \times 1.000 \times 1.38 = 14{,}802 \text{ RMB}. \tag{5.4}$$

The T2S1 group DeepSurv rate mean is 27,198 RMB—the case individual's Tweedie baseline premium (7,155 RMB) is below the group mean, but the DeepSurv risk multiplier (1.28) and the sojourn-time factor (1.171) adjust it upward, so that the final gross premium of 14,802 RMB reflects the individual's moderately high composite risk level. For comparison, if the traditional EM model's T2S1 factor of 2.18 times (based on the

T2S0 baseline of 3,534 RMB) were used, the gross premium would be only approximately 10,630 RMB, underestimating the DeepSurv-justified level by about 28%—the distortion of the EM factor is clearly amplified in individual pricing, highlighting the pricing correction value of the deep model.

Table 9 presents the SHAP premium attribution table for this case. The base value $\phi_0 = 0.50$ represents the population-average risk level. Age 65 contributes the most, with a SHAP value of +0.74 (38.1%); DX_T2S1 contributes +0.43 (22.2%); heart failure COMB_I50 contributes +0.33 (17.0%); chronic ischemic heart disease COMB_I25 contributes +0.32 (16.5%); and male sex contributes +0.12 (6.2%). The SHAP values of all features sum to $f_\theta = 2.44$, corresponding to a gross premium of 14,802 RMB.

Table 9 SHAP premium attribution table for Case A

Item	SHAP value	Contribution share
Base value ϕ_0	0.50	—
Age (65 years)	+0.74	38.1%
DX_T2S1	+0.43	22.2%
COMB_I50 (heart failure)	+0.33	17.0%
COMB_I25 (chronic ischemic heart disease)	+0.32	16.5%
Sex (male)	+0.12	6.2%
Total f_θ	2.44	100%

The significance of this attribution table transcends the numbers themselves. It makes an ostensibly cold, high premium amount traceable and communicable: the policyholder can clearly see why their premium is at this level, which comorbidity contributed how much, and how much they could save by improving which condition. For regulators, it provides a basis for reviewing rate fairness; for insurers, it translates the deep model's output into business language. Interpretability is thus elevated from a technical requirement to a vehicle for pricing ethics—this is the deeper motivation behind our insistence on SHAP attribution.

6 Conclusion and Outlook

Taking type 2 diabetes mellitus as a case study, this paper has constructed a substandard-life insurance rate-making framework centered on the DeepSurv deep survival model, with Elastic Net regularization as the shrinkage mechanism, LSTM temporal prediction as an independent comparative benchmark, SHAP as the interpretability outlet, and Bootstrap as the robustness guarantee, completing the full chain from risk modeling to rate implementation.

The main conclusions are fivefold. First, traditional methods exhibit systematic deficiencies: the EM model, failing to control for age confounding, produces a rate inversion between T2S1 (118.22%) and T2S2 (88.87%), with Bootstrap confirming the stability of this inversion in the sampling distribution; the GLM two-part approach, through covariate adjustment, repairs the inversion and decomposes the structural divergence of T2S2 "high frequency" (Poisson multiplier 1.89) and T2S1 "high severity" (Gamma multiplier

1.65), but after incorporating province and year controls, the age coefficient in the Gamma severity model, while corrected from -0.0002 to $+0.0028$ (the directional bias is repaired), remains far below the Poisson frequency coefficient of 0.031, indicating the limited capacity of the linear-additive structure to capture age effects; the Tweedie, with $p = 1.90$, achieves compound modeling on the full data (RMSE = 33,823 RMB) and appropriately allocates the age effect across the frequency and severity dimensions, but its capture of nonlinear interactions among comorbidities remains constrained by the additive structure.

Second, DeepSurv captures the high-dimensional interactions of comorbidities through a nonlinear risk surface, and the Elastic Net regularization achieves feature selection and model simplification without sacrificing discriminative power. On the full 1,840,774 records, DeepSurv attains a C-index of 0.695. It should be emphasized that the Tweedie/GLM AUC for binary discrimination of excess events reaches 0.715, but the two metrics are not directly comparable—AUC measures binary classification power for whether out-of-pocket expenses exceed a 10,000 RMB threshold, while the C-index measures concordance in survival time ranking. The core value of DeepSurv lies not in a dominant single discrimination metric, but in its nonlinear risk surface that can automatically encode high-dimensional interactions among comorbidities and enable premium attribution through SHAP—capabilities beyond the reach of linear models. The comparison experiment of three regularization schemes reveals that while all three groups show 0% sparsity (i.e., no feature is driven to exact zero), the differences in Gini coefficients (Elastic Net 0.506 vs. pure L2 0.488 vs. pure L1 0.493) indicate that the Elastic Net achieves the optimal balance in weight stratification—assigning similar weights to the collinear comorbidity pair I50 and I25, avoiding the arbitrary selection of pure Lasso within collinear groups and the inability of pure Ridge to perform variable selection. Bootstrap resampling ($B = 50$) confirms a C-index standard error of only 0.0031, with a 95% CI of $[0.684, 0.695]$; the Tweedie AUC confidence interval is $[0.715, 0.716]$. Although the two intervals do not numerically overlap, this does not constitute a statistical verdict of model superiority, but rather reflects two different facets of model capability captured by the two metrics.

Third, the LSTM temporal prediction, as an independent comparative experiment, verifies the incremental value of temporal information: the temporal R^2 of 0.391 outperforms the GLM cross-section (improvement of approximately 99%), but does not surpass the lagged-mean benchmark (0.455). Subgroup analysis reveals that temporal information has the greatest predictive value in the T2S1 (moderate risk, variable trajectory) population ($R^2 = 0.552$), providing a precise application locus for the future embedding of LSTM within DeepSurv.

Fourth, SHAP gradient approximation on the full 1.84 million sample identifies age (0.269), DX_T2S2 (0.129), chronic ischemic heart disease COMB_I25 (0.123), DX_T2S1 (0.107), and heart failure COMB_I50 (0.107) as the top five risk drivers, with cardiovascular comorbidities approaching or even exceeding diabetes staging itself in attribution weight. It is noteworthy that the SHAP value of I25 exceeds that of I50, contrasting with the Tweedie model where the marginal coefficient of I50 is larger—the deep network, through hidden-layer interactions, re-encodes the synergistic effect of I25 as an upstream cause into the risk score, a nonlinear structure that the linear-additive model fails to fully characterize.

Fifth, the Bootstrap robustness checks cover all key metrics, including model discriminative power, rate levels, loss ratios, loss probabilities, and EM factors. All metrics exhibit standard errors within acceptable ranges (e.g., gross premium SE is only 0.4%–

0.5% of the mean), and the 95% confidence intervals are compact, confirming that the research conclusions are reproducible under repeated sampling. This robustness guarantee is particularly critical for the application of deep models in actuarial regulatory contexts.

The applied value of the method is concretely demonstrated in the complete pricing case: the high-risk Case A yields a gross premium of 14,802 RMB, whereas the EM model would give only approximately 10,630 RMB, underestimating by about 28%. Coupled with the sojourn-time factor $A(\tau)$ and the health-management factor $H(r)$, the framework can both reflect disease-course dynamics and endogenize the positive incentive of health management into the rate—a T2S2 patient fully participating in health management could receive a premium discount of up to 28.2%.

Several directions merit further exploration. At the data level, this study uses inpatient claims as the primary endpoint; future work could incorporate outpatient and medication data to more completely characterize the disease course. Explicit modeling of competing risks (rather than treating them as censored) also warrants deeper investigation. At the methodological level, the joint tuning of α and λ in the Elastic Net regularization (e.g., via cross-validation on a two-dimensional grid) could be further refined; the joint training of DeepSurv and LSTM holds promise for further unlocking the value of temporal information; and the fusion of semi-Markov multi-state processes [15] with deep survival models could more rigorously characterize state transitions. Bootstrap is currently performed only at the evaluation level; if computational resources permit, future work could attempt resampling at the training level to assess the uncertainty of model parameters. At the application level, the regulatory compliance of rate factors and the privacy–utility trade-off in cross-institutional federated training are practical issues that require ongoing refinement in the implementation process. The technical evolution of substandard-life insurance ultimately serves a simple goal: to ensure that every impaired-life applicant can obtain fair, transparent, and affordable risk protection.

References

[1] Cox D R. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 1972, 34(2): 187–220.

[2] Katzman J L, Shaham U, Cloninger A, et al. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Medical Research Methodology*, 2018, 18(1): 24.

[3] Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Computation*, 1997, 9(8): 1735–1780.

[4] Lundberg S M, Lee S I. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 2017, 30: 4765–4774.

[5] Shapley L S. A value for n-person games. In: *Contributions to the Theory of Games II*. Princeton: Princeton University Press, 1953: 307–317.

[6] Tweedie M C K. An index which distinguishes between some important exponential families. In: *Statistics: Applications and New Directions*. Calcutta: Indian Statistical Institute, 1984: 579–604.

[7] Smyth G K, Jørgensen B. Fitting Tweedie's compound Poisson model to insurance claims data: dispersion modelling. *ASTIN Bulletin*, 2002, 32(1): 143–157.

[8] Nelder J A, Wedderburn R W M. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 1972, 135(3): 370–384.

[9] Schoenfeld D. Partial residuals for the proportional hazards regression model. *Biometrika*, 1982, 69(1): 239–241.

[10] Kaplan E L, Meier P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 1958, 53(282): 457–481.

[11] Nelson W. Hazard plotting for incomplete failure data. *Journal of Quality Technology*, 1969, 1(1): 27–52.

[12] van der Vaart A W. *Asymptotic Statistics*. Cambridge: Cambridge University Press, 1998.

[13] Cybenko G. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 1989, 2(4): 303–314.

[14] Fleming W H, Soner H M. *Controlled Markov Processes and Viscosity Solutions*. 2nd ed. New York: Springer, 2006.

[15] Janssen J, Manca R. *Applied Semi-Markov Processes*. New York: Springer, 2006.

[16] Lemaire J. *Bonus-Malus Systems in Automobile Insurance*. Boston: Kluwer Academic Publishers, 1995.

[17] UK Prospective Diabetes Study (UKPDS) Group. Intensive blood-glucose control with sulphonylureas or insulin compared with conventional treatment and risk of complications in patients with type 2 diabetes (UKPDS 33). *The Lancet*, 1998, 352(9131): 837–853.

[18] International Diabetes Federation. *IDF Diabetes Atlas*. 10th ed. Brussels: International Diabetes Federation, 2021.

[19] Qiu C J, Li Y H, Tan X Y, Qian L Y. Pricing critical illness insurance based on a continuous-time three-state Markov model. *Statistical Research*, 2023, 40(2): 88–102. (In Chinese)

[20] Bladt M. Phase-type distributions in life and survival modelling: recent advances. *Scandinavian Actuarial Journal*, 2025, 2025(2): 105–130.

[21] Boucher J P, Turcotte R. Modelling insurance claim durations: a survival analysis perspective. *North American Actuarial Journal*, 2024, 28(1): 55–78.

[22] Gu J, Yang Y. Deep learning for loss reserving and survival-based pricing. *Insurance: Mathematics and Economics*, 2024, 116: 210–228.

[23] Zeng H, et al. A DeepSurv model for prognosis prediction in gastrointestinal cancer with a concordance index of 0.825. *Frontiers in Oncology*, 2024, 14: 1–12.

[24] Netayawijit P, et al. Interpretable machine learning for diabetes risk prediction using SHAP. *Health Informatics Journal*, 2025, 31(1): 1–18.

[25] Brackenridge R D C, Elder W J. *Medical Selection of Life Risks*. 4th ed. London: Macmillan Press, 1998.

[26] Zou H, Hastie T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2005, 67(2): 301–320.

[27] Yang Y, Schulhofer-Wohl S, Fu W J, Land K C. The intrinsic estimator for age-period-cohort analysis: what it is and how to use it. *American Journal of Sociology*, 2008, 113(6): 1697–1736.