International Journal of Computational and Experimental Science and ENgineering
(IJCESEN)

Vol. 12-No.1 (2026) pp. 1-7
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Profitability Impact of AI Adoption in Annuity Pricing

Padmaja Dhanekulla*

Akkodis Group/Principal Software Engineer
* **Corresponding Author Email:** pdhanekulla@gmail.com- **ORCID:** 0000-0002-5947-3350

Article Info:

DOI: 10.22399/ijcesen.4599
Received : 05 November 2025
Revised : 27 December 2025
Accepted : 29 December 2025

Keywords

Artificial Intelligence Pricing Models,
Annuity Product Optimization,
Machine Learning Actuarial Science,
Spread Margin Enhancement,
Stochastic Mortality Forecasting,
Reinforcement Learning Hedging

Abstract:

The insurance enterprise faces growing challenges in optimizing annuity pricing strategies under changing market situations and growing opposition to its product portfolios. Conventional actuarial models have massive deficiencies, as they depend upon historical data patterns and linear assumptions that cannot capture complex, nonlinear relationships among risk factors and profitability results. The usage of machine learning algorithms in annuity pricing frameworks truly addresses core shortcomings in conventional methodologies, with particular regard to mortality risk assessment, sensitivity analysis of interest rates, and policyholder behavior prediction. Gradient boosting decision trees, neural network architectures, and ensemble methods applied to fixed and indexed annuity products achieve superior predictive accuracy compared to generalized linear models commonly found in traditional actuarial practice. This implementation framework integrates machine learning predictive capabilities with established actuarial practices to ensure regulatory compliance while maintaining mathematical soundness. Deep learning approaches to mortality forecasting, especially Long Short-Term Memory networks, transcend restrictions of classical Lee-Carter models by allowing temporal dependencies and nonlinear patterns to be captured, characteristic of modern mortality experiences. Reinforcement learning applications to derivative hedging strategies optimize dynamic rebalancing decisions for indexed annuity products with embedded options. Quantitative comparison among stochastic mortality models across heterogeneous populations reveals performance differences conditional upon demographic characteristics and projection horizons. Spread margin enhancement through predictive analytics allows a more sophisticated crediting rate determination and policyholder retention strategy. ROI considerations include not only direct improvements in profitability but also indirect operational efficiency gains, balanced against infrastructure and personnel investment requirements.

1. Introduction

Annuity pricing is a major profit center for insurance carriers, and even a small, marginal benefit in spread margins results in sizeable advantages and higher profitability. The basic challenge remains precise forecasting of mortality rates, persistency patterns, and investment returns across various product lines with effective management of capital requirements and regulatory constraints. Traditional pricing approaches use static actuarial tables and point deterministic assumptions that do not accurately capture the dynamic nature of underlying risk factors. The classical Lee-Carter model, notwithstanding its widespread application in mortality forecasting since the early 1990s, uses singular value decomposition to extract temporal mortality patterns and age-specific sensitivities but fails to capture nonlinear dependencies and cohort-specific effects that typify modern mortality experiences. Recent breakthroughs in deep learning architectures have illustrated the ability to extend traditional stochastic mortality models using recurrent neural network frameworks, particularly Long Short-Term Memory networks, which process sequential mortality data while maintaining the interpretability requirements necessary for actuarial validation and regulatory acceptance [1]. The increasingly available sources of alternative data, such as wearable device metrics, socioeconomic indicators, and behavioral analytics, are creating an opportunity for more fine-grained risk segmentation than what traditional models can exploit meaningfully.

Machine learning approaches provide the ability to process high-dimensional feature spaces and detect non-linear patterns driving the pricing outcomes. Application of ensemble methods and multivariate regression techniques in modelling insurance claims has shown that these methods clearly outperform univariate techniques with separate models for frequency and severity. Research into motor third-party liability insurance has demonstrated the remarkable ability of generalized linear models, once expanded by a carefully designed set of interaction terms and hierarchical structure, to capture complex relationships between policyholder characteristics, vehicle attributes, and claims outcomes, though at the cost of computational intensity when model complexity increases [2]. However, widespread adoption within the insurance industry is slowed by increased scrutiny from regulators over model interpretability, growing concern over algorithmic bias, and the need to integrate complex artificial intelligence systems into legacy actuarial workflows. A central tension yet remains the trade-off between predictive accuracy and explainability, whereby black-box algorithms superior in performance find resistance in regulatory frameworks centered on transparent, auditable methodologies.

This research addresses these barriers by developing a hybrid framework that combines machine learning predictive power with actuarial principles and thus assures both accuracy and regulatory compliance. The integration of neural network architectures with established stochastic models provides a pathway to enhance forecasting capabilities while maintaining the theoretical foundations that underpin risk-based pricing. This investigation focuses specifically on spread margin optimization, wherein the profitability of products is determined by the difference between credited rates and investment returns earned. Utilizing techniques proven highly effective in both mortality forecasting and claims modelling applications, the framework extends such methodologies to the unique challenges posed within annuity pricing—including the multi-decade time horizons, the embedded optionality in indexed products, and the complex interplay between asset-liability management and competitive positioning considerations.

2. Machine Learning Architecture for Pricing Optimization

The application of AI-driven pricing models requires the careful selection of algorithms that balance predictive performance with requirements for computational efficiency and interpretability. Gradient boosting decision trees provide a strong foundation for dealing with the mixed data types present within insurance datasets, which can include continuous variables such as age and premium amount, through to categorical features relating to underwriting class and distribution channel. The iterative nature of gradient boosting allows the model to progressively refine predictions by focusing on residual errors, making it particularly well-suited to capturing subtle patterns in policyholder behavior and mortality risk. Various trials of recurrent neural network architectures, particularly Long Short-Term Memory networks and Gated Recurrent Units, have shown considerable promise in mortality rate forecasting tasks via their ability to model temporal dependencies in sequential data. Comparative tests conducted against the benchmark Lee-Carter model demonstrate that recurrent architectures are capable of capturing non-linear patterns and sudden shocks in mortality trends that traditional stochastic models cannot accommodate, although performance gains vary substantially depending on the length of available historical data series and the degree of structural breaks present in mortality experiences [3].

Neural network architectures supplement these with their inherent capability for hierarchical feature representation of raw data in an automated way. Multi-layer perceptrons with carefully designed activation functions can model complex interactions between pricing variables that would require extensive manual feature engineering in traditional approaches. The incorporation of neural network layers within more classical statistical frameworks is a most promising direction for actuarial applications, as these combine the interpretability advantages from established methods with flexible function approximation capabilities afforded by deep learning. Hybrid architectures embedding neural network transformations within logistic regression frameworks enable the model to learn non-linear feature representations while maintaining the coefficient-based structure familiar to both actuarial practitioners and regulatory auditors. Studies of this class of hybrid have shown that incorporation of hidden layers between raw inputs and logistic output significantly enhances classification accuracy in insurance risk classification tasks, with the neural components of the model effectively conducting automatic feature engineering otherwise requiring domain expertise and iterative model refinement [4].

Dropout regularization and batch normalization prevent overfitting while sustaining generalization performance across diverse market segments. Ensemble methods that combine multiple model types leverage their respective strengths, with mechanisms for voting or stacked architectures producing more robust predictions than any single algorithm. The approach of stacking utilizes a meta-learner in order to integrate predictions emanating from diverse base models efficiently; it learns the optimal weighting scheme concerning different algorithms on the basis of performance across the segments of the feature space. Recurrent architectures have been really effective in instances where forecasting mortality requires the model to identify regime changes or structural breaks in underlying trends, as their intrinsic gating mechanisms in the Long Short-Term Memory networks allow the selective retention of relevant historical information while discarding outdated patterns that do not reflect the current mortality dynamics anymore [3].

Feature engineering is a necessary step that converts domain knowledge into model inputs, which in turn drive predictions. Temporal features that capture policy duration, seasonal patterns, and economic cycle effects provide context for understanding persistence behavior. Interaction terms between age bands and product types will allow the model to learn segment-specific pricing sensitivities. Including lagged variables representing the historical actions of policyholders will enable the system to pick out patterns indicative of future behavior, such as surrender risk or additional premium contributions. The neural network components in hybrid models can learn feature interactions directly through non-linear transformations across hidden layers. This reduces manual efforts compared to traditional feature engineering while potentially uncovering novel risk factors not previously identified with conventional actuarial analysis. The challenge, however, is model transparency; the predictive power gain comes with reduced interpretability compared with purely linear frameworks [4].

3. Improving Spread Margin Using Predictive Analytics

The optimization of spread margins necessitates a precise forecast of the earned rate on invested assets and the credited rate promised to the policyholders. Machine learning models enhance this process through better predictive analytics on the returns of investment portfolios under different economic scenarios. Time series forecasting algorithms, such as Long Short-Term Memory networks and temporal convolutional networks, process historical yield curves and macroeconomic indicators to project future interest rate environments. Long Short-Term Memory architectures have been applied to portfolio optimization tasks, revealing high potential in quantitative finance applications, where the core recurrent network processes sequential market data in order to detect patterns that are predictive of subsequent asset returns. Its ability to maintain temporal dependencies over long periods makes it possible to put fundamental economic indicators, technical price patterns, and market sentiment variables within a unified forecasting framework. Portfolio allocation strategies based on the Long Short-Term Memory model predictions are leveraging its ability to learn how to update weight parameters via the process of backpropagation through time to learn which historical features provide the most information for forecasting returns during various regimes and conditions of volatility [5].

These projections are used to inform dynamic crediting rate strategies that maintain competitive positioning, while protecting profitability. Predictive analytics enables sophisticated approaches to bonus rate determination in indexed annuity products. By modelling the complex relationship between participation rates, cap rates, and underlying index performance, artificial intelligence systems identify optimal combinations of parameters that maximize customer value perception while controlling option costs. Correlation structures across multiple indices, volatility surface dynamics, and policyholder election patterns are considered to inform product design decisions. The enhanced Monte Carlo simulation with machine learning predictions creates probability distributions of profitability outcomes under a variety of market scenarios, which better inform risk management.

Policyholder behavior modelling represents another crucial element of spread margin optimization. Neural networks trained on the historical surrender data may highlight early warnings of the termination of policies and, therefore, proactive retention strategies, coupled with better cash flow forecasting. The essential challenge to annuity pricing is the projection of future mortality across the long projection periods typical of such long-duration contracts, where small errors in the assumptions about mortality rates amplify significantly over multi-decade projection horizons. Parametric mortality models, such as several extensions and variants of the Lee-Carter framework, provide formal approaches to extrapolating historical mortality trends forward in time while retaining mathematical tractability and actuarial interpretability. Research into methods for projecting mortality for immediate annuitants and life office pensioners has demonstrated that different parametric specifications fit to different extents depending on particular population characteristics and the historical period considered. The choice of appropriate model structures involves careful trade-offs between model complexity and forecast stability, with highly elaborate specifications fitting the historical data more closely but often generating unreliable long-term projections due to instability of parameter estimates [6].

These models integrate the impact of competitive product offerings and current market rates, together with policy-specific features, to deliver individual risk scores. This gives a granular insight that can allow dynamic pricing and targeted product recommendations to enhance portfolio profitability overall. The parametric approach to mortality forecasting provides a natural framework for incorporating cohort effects and period-specific influences on age-specific death rates, making it possible to achieve more refined annuity product pricing that needs to be relevant for a wide range of demographic segments and changing longevity patterns [6].

4. Implementation Framework for Fixed and Indexed Annuities

Deployment of the AI-driven pricing system in fixed annuity products requires the integration of actuarial reserving methodologies and financial reporting frameworks. The architecture for such an implementation utilizes a microservices approach wherein machine learning models are built as individual nodes to interface with the core policy administration systems through well-defined application programming interfaces. This allows for stability in legacy systems, enabling rapid iteration and enhancement of predictive models. All scoring capabilities for quote generation are real-time; model predictions feed directly into pricing engines that determine final premium rates and credited rates. Architecturally, the considerations for deployment in production insurance environments go beyond model accuracy to include latency requirements, fault tolerance, and integration with existing data infrastructure.

For indexed annuity products, this complexity escalates since the embedded derivative features demand sophisticated hedging strategies. This process is supported by the machine learning models that forecast the option cost and optimal hedge ratios under various market conditions. Thus, the system keeps track of real-time market data streams such as equity index levels, implied volatility surfaces, and interest rate term structures to develop continuous pricing recommendations. Hedging decisions can be optimized through reinforcement learning algorithms learning from historical profit and loss outcomes, incorporating methods that can make trial-and-error interactions with simulated market environments towards progressive strategy improvement. Deep reinforcement learning approaches to derivative hedging have illustrated substantial advantages over traditional delta-gamma hedging frameworks by learning the optimal rebalancing policy that incorporates transaction costs, market impact, and incomplete market conditions. The application of deep deterministic policy gradient algorithms and proximal policy optimization methods allows agents to build hedging strategies that minimize variance while controlling the trade-off between hedging effectiveness and trading costs. The neural network policy maps observable market states such as underlying asset prices, historical volatility measures, and current portfolio positions to continuous hedging actions that adapt dynamically to the changing market conditions rather than static rules obtained using Greeks-based derivations from idealized Black-Scholes assumptions [7].

Model governance and validation procedures guarantee that AI-based pricing systems remain accurate and uphold regulatory requirements. Backtesting frameworks pit predicted outcomes against experience in a wide range of time periods and market cycles. Sensitivity analysis looks into model behavior under extreme scenarios to find possible weaknesses. The validation of stochastic mortality models requires extensive comparisons across multiple methodological frameworks so as to assess relative performance for different population characteristics and historical periods. Quantitative evaluations of mortality projection models on Italian population data have indicated that different parametric specifications yield quite varying degrees of forecast accuracy as one varies the age cohort, projection horizons, and the specific demographic structure of the underlying population. The Lee-Carter model and its variants, including cohort adjustment and multiple population settings, present distinct profiles of fit when applied to mortality data with non-smooth historical trends or structural breaks in the improvement rates. Model selection criteria need to balance the historical goodness-of-fit against the stability of the projection, since model specifications that provide the best fit to past mortality may produce implausible long-run forecasts if historical patterns prove to be non-stationary [8].

Explainability techniques, such as Shapley Additive Explanations values and partial dependence plots, allow detailed insight into model decision-making processes, both regulatory and trust-building in actuarial teams. Continuous monitoring detects model drift and triggers retraining procedures upon the occurrence of performance degradation. This relative assessment of the mortality models across national populations evidences that regular recalibration is important when new data from experience emerges; indeed, there is evidence to believe that mortality improvement patterns tend to vary greatly over both space and time [8].

5. Return on Investment Considerations and Performance Metrics

Assessing the economic impact of artificial intelligence adoption requires comprehensive frameworks that capture both direct and indirect advantages. Direct profitability improvements manifest through enhanced spread margins resulting from more accurate pricing and improved persistency rates achieved through better customer segmentation. The measurement of these effects requires careful attribution analysis to isolate the impact of AI-driven changes from confounding factors such as market movements and competitive dynamics. Cohort analysis comparing policies priced using traditional methods versus AI-enhanced approaches provides evidence of performance differences whilst controlling for temporal effects. The proliferation of machine learning applications across financial services has demonstrated diverse implementation patterns, with supervised learning techniques dominating credit risk assessment and insurance underwriting tasks, whilst unsupervised methods find application in fraud detection and customer segmentation challenges. The landscape of machine learning deployment in finance reveals that gradient boosting frameworks and random forest algorithms remain prevalent due to their robust performance on structured data typical of financial transactions, though neural network architectures gain traction for unstructured data processing, including text analysis of policy documents and sentiment extraction from customer communications. The practical implementation of these technologies faces obstacles, including data quality concerns, regulatory compliance requirements, and the challenge of integrating novel methodologies within established risk management frameworks that emphasize interpretability and auditability [9].

Indirect benefits include operational efficiency gains from automated pricing processes and reduced actuarial workload. The quantification of these advantages considers time savings in pricing studies, reduced error rates in manual calculations, and faster time-to-market for new product launches. Infrastructure costs associated with cloud computing resources, data storage, and specialized personnel represent significant investment requirements that must be balanced against projected benefits. The analysis employs discounted cash flow methodologies to evaluate long-term return on investment horizons appropriate for insurance business models. Financial institutions report that successful machine learning adoption requires substantial organizational change management alongside technical implementation, given that actuarial teams need to develop proficiency in model validation techniques adapted to non-linear algorithms that lack closed-form solutions [9].

Risk-adjusted performance metrics give a more comprehensive view of the outcomes from implementing artificial intelligence by accounting for reduced adverse selection and improved capital efficiency. It creates value through better return on allocated capital ratios resulting from the decrease in required economic capital due to better risk segmentation. The final component, scenario testing, assesses system performance under various stress conditions, including extreme market disruptions and unexpected mortality events, to ensure the resilience of the AI-driven pricing strategy. Learning curves incorporated into the framework mirror improving model performance as training data sets expand and algorithmic refinements accumulate over time. Insurance risk classification often involves ordinal categories showing different severity levels or rating classes that naturally have ordering relationships. Deep neural network architectures for ordinal regression provide clear advantages over the standard approach to classification, explicitly incorporating rank consistency constraints that ensure predicted probabilities respect the natural ordering of the risk categories. Such architectures use conditional probability frameworks that decompose ordinal prediction into sequences of binary classification, which allows the network to learn threshold parameters that delineate adjacent risk class boundaries while preserving monotonicity properties so crucial for actuarial credibility.

Table 1. *Machine Learning Algorithms for Annuity Pricing Applications [3, 4].*

Algorithm	Application	Key Capabilities
Gradient Boosting Trees	Risk classification, mortality prediction	Mixed data handling, iterative error reduction, and non-linear pattern capture
LSTM Networks	Temporal mortality forecasting	Long-range dependencies, sequential processing, regime change detection
Multi-Layer Perceptrons	Feature representation	Automatic interaction detection, hierarchical transformations
Hybrid Neural-Logistic	Risk categorisation	Non-linear learning with coefficient-based interpretability
Ensemble Stacking	Multi-model integration	Meta-learner optimisation, complementary strength leverage

Table 2. *Predictive Analytics for Spread Margin Optimisation [5, 6]*

Component	Target	Methodology	Application
Interest Rate Projection	Investment returns	LSTM with yield curves	Dynamic crediting strategies
Option Cost Estimation	Derivative pricing	Neural networks with volatility surfaces	Cap and participation rates
Lapse Prediction	Surrender probability	Gradient boosting with market factors	Retention and cash flow forecasting
Mortality Forecasting	Long-term death rates	Parametric models with cohorts	Annuitant pricing and reserves
Portfolio Simulation	Profitability distribution	ML-enhanced Monte Carlo	Risk management validation

Table 3. *AI-Driven Pricing System Implementation [7, 8].*

Component	Technical Approach	Purpose
Microservices Architecture	API-integrated ML modules	Real-time quote generation
Deep RL Hedging	Policy gradient algorithms	Optimal derivative rebalancing
Backtesting Framework	Multi-cycle validation	Performance monitoring
Explainability Tools	SHAP values, dependence plots	Regulatory approval support
Drift Detection	Automated monitoring	Triggers model retraining
Mortality Validation	Parametric comparison	Projection stability assessment

Table 4. *Return on Investment Assessment [9, 10].*

Benefit	Measurement	Value Driver
Spread Margin Enhancement	Cohort analysis	Improved pricing accuracy
Operational Efficiency	Time savings quantification	Process automation
Capital Efficiency	Economic capital reduction	Better risk segmentation
Adverse Selection Reduction	Risk-adjusted metrics	Ordinal regression classification
Infrastructure Investment	Discounted cash flow	Cloud computing capabilities
Learning Curve Gains	Accuracy tracking	Dataset expansion and refinement

6. Conclusions

The transformation of annuity pricing practice by means of AI represents a paradigm shift from classical deterministic frameworks to adaptive systems that learn continuously from emerging data streams. Once set up within appropriate actuarial structures, machine learning algorithms produce measurable improvements in pricing accuracy and profitability for fixed and indexed product lines. Gradient boosting and neural network architectures, which were discussed throughout the article, successfully captured nonlinear relationships between risk factors that are poorly represented by conventional generalized linear specifications. Advanced techniques applied to spread margin optimization create competitive advantages through improved interest rate forecasting, policyholder behavior prediction, and dynamic crediting rate strategies responsive to market conditions. A microservices implementation framework addresses critical regulatory compliance challenges, model interpretability requirements, and legacy system integration constraints. Governance procedures ensure that models remain accurate over time via comprehensive validation methodologies, backtesting protocols, and transparent documentation practices that facilitate regulatory acceptance. Applications of deep reinforcement learning to derivative hedging demonstrated substantial gains over traditional delta-gamma approaches through the learning of optimal rebalancing policies that account for transaction costs and market frictions. Quantitative comparisons among stochastic mortality models revealed performance dependencies on population characteristics and projection horizons, highlighting the importance of regular recalibration as experience data accumulates. Return on investment analyses confirmed that benefits accrue beyond direct profitability to encompass operational efficiencies and enhanced risk management capabilities. Infrastructure and personnel investments required for successful implementation generate value through adverse selection reduction and improved capital efficiency. Future work should examine attention mechanisms for time series forecasting, graph neural networks for portfolio dependency modelling, and causal inference methodologies that strengthen the attribution of profitability improvements to specific interventions. Regulatory frameworks continue to evolve to accommodate algorithmic pricing, expanding opportunities for AI-enhanced actuarial practice and driving sustained innovation in methodology and implementation strategies.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] Andrea Nigri et al., "A Deep Learning Integrated Lee–Carter Model," MDPI, 2019. [Online]. Available: https://www.mdpi.com/2227-9091/7/1/33

[2] Aivars Spilbergs et al., "MULTIVARIATE MODELLING OF MOTOR THIRD PARTY LIABILITYINSURANCE CLAIMS, "European Journal of Business Science and Technology, 2022

[3] G´abor Petneh´azi and J´ozsef G´all, "Mortality rate forecasting: can recurrent neural networks beat the Lee-Carter model," arXiv, 2019. [Online]. Available: https://arxiv.org/pdf/1909.05501

[4] George Tzougas and Konstantin Kutzkov, "Enhancing Logistic Regression Using Neural Networks for Classification in Actuarial Learning," MDPI, 2023. [Online]. Available: https://www.mdpi.com/1999-4893/16/2/99

[5] Van-Dai Ta et al., "Portfolio Optimization-Based Stock Prediction Using Long-Short Term Memory Network in Quantitative Trading," MDPI, 2020. [Online]. Available: https://www.mdpi.com/2076-3417/10/2/437

[6] Terry Z. Sithole et al., "An investigation into parametric models for mortality projections, with applications to immediate annuitants' and life office pensioners' data," ResearchGate, 2000. [Online]. Available: https://www.researchgate.net/profile/Steven-Haberman/publication/222837054

[7] Jay Cao et al., "Deep Hedging of Derivatives Using Reinforcement Learning," Journal of Financial Data Science, 2021. [Online]. Available: https://arxiv.org/pdf/2103.16409

[8] M.F. Carfora et al., "A quantitative comparison of stochastic mortality models on Italian population data," ResearchGate, 2017. [Online]. Available: https://www.researchgate.net/profile/Maria-Francesca-Carfora/publication/315322881

[9] Daniel Hoang and Kevin Wiegratz, "Machine learning methods in finance: Recent applications and prospects," Wiley, 2023. [Online]. Available: https://onlinelibrary.wiley.com/doi/pdfdirect/10.1111/eufm.12408

[10] Xintong Shi et al., "Deep Neural Networks for Rank-Consistent Ordinal Regression Based on Conditional Probabilities," arXiv, 2023. [Online]. Available: https://arxiv.org/pdf/2111.08851