Science Discovery Mathematics
2026, Vol. 1, No. 2, pp. 48–53
https://doi.org/10.11648/j.sdmath.20260102.11

Research Article

Uncertainty-Calibrated Distributionally Robust Pricing for Insurance: From AI-Based RiskPrediction to Robust and Non-Conservative Premiums

Majid Ghorbani*

Department of Mathematics, Azad University, Tehran, Iran

Abstract

The increasing use of machine learning and artificial intelligence in insurance pricing has created new opportunities for modeling complex relationships between policyholder characteristics and insurance losses. Unlike conventional actuarial models, modern probabilistic machine-learning approaches can provide flexible predictive distributions and can potentially quantify different forms of predictive uncertainty. Nevertheless, the availability of sophisticated predictive distributions does not by itself solve the problem of decision-making under uncertainty. A pricing model must determine how uncertainty in the estimated loss distribution should affect the premium charged to an insured risk. Distributionally robust optimization provides a natural framework for addressing uncertainty in probability distributions by replacing a single reference distribution with an ambiguity set containing plausible alternatives. Although distributionally robust methods and their applications to insurance are already well established, an important methodological issue remains concerning the connection between predictive uncertainty generated by modern AI models and the construction of ambiguity sets for robust insurance pricing. In particular, ambiguity sets that are excessively small may fail to provide adequate robustness, whereas overly large sets can produce excessively conservative premiums. This paper develops a conceptual mathematical framework for uncertainty-calibrated distributionally robust insurance pricing. The framework connects four components: AI-based probabilistic risk prediction, predictive uncertainty quantification, statistically calibrated ambiguity-set construction, and robust premium determination. Particular attention is given to the relationship between predictive calibration, ambiguity-set size, robustness guarantees, and premium conservatism. The paper does not claim that distributionally robust insurance pricing or machine-learning-based insurance pricing is novel in isolation. Instead, it identifies a focused research direction concerning the mathematical calibration of distributional ambiguity from AI-derived predictive uncertainty and the characterization of its consequences for robust insurance premiums. The proposed framework provides a basis for future theoretical work on stability, sensitivity, convergence, and conservatism bounds in uncertainty-aware insurance pricing.

Keywords

Insurance Pricing, Machine Learning, Artificial Intelligence, Predictive Uncertainty, Distributionally Robust, Optimization

*Correspondence: Majid Ghorbani (ghorbani10100@gmail.com)

Received: 20 August 2026; **Accepted:** 31 August 2026; **Published:** 18 September 2026

1. Introduction

Insurance pricing is fundamentally a problem of decision-making under uncertainty. An insurer must estimate the expected cost and risk associated with a policy or portfolio and then determine an appropriate premium that is sufficiently adequate to cover future losses while remaining economically and competitively meaningful. For many decades, actuarial pricing has relied on statistical models designed to describe claim frequency, claim severity, aggregate losses, and relationships between risk characteristics and expected claims.

The increasing availability of high-dimensional insurance data has substantially expanded the modeling possibilities. Policy characteristics, historical claims, behavioral information, telematics, geographic variables, environmental conditions, and other sources of information can contain complex nonlinear relationships that may not be easily represented through traditional parametric structures. Machine-learning methods have consequently become increasingly relevant to actuarial science and insurance pricing. A comprehensive review by Blier-Wong, Cossette, Lamontagne, and Marceau documented the growing role of machine learning in property and casualty insurance and highlighted the movement from conventional generalized linear models toward more flexible predictive methods. [1]

The use of machine learning, however, creates an important distinction between prediction and decision-making. A model may provide accurate predictions while producing outputs that are poorly calibrated for pricing purposes. In insurance, this issue is particularly important because premiums must satisfy actuarial balance properties and should remain consistent with observed losses. Denuit, Charpentier, and Trufin demonstrated that machine-learning-based insurance pricing can exhibit balance problems and proposed auto-calibration as a mechanism for improving local balance properties. [2]

A second problem arises from the uncertainty of the predictive distribution itself. Even if a machine-learning model produces a probability distribution for future losses, that distribution has been estimated from finite data. Its parameters, functional form, calibration, and generalization to future observations may therefore be uncertain. A pricing procedure based entirely on a single estimated distribution can consequently underestimate distributional uncertainty.

Distributionally robust optimization offers a different perspective. Rather than assuming that a single estimated probability distribution represents the true loss-generating process, DRO considers a family of plausible distributions and evaluates decisions against adverse members of that family. The ambiguity set is therefore a central object in the formulation. Modern DRO theory includes statistically motivated ambiguity sets, dual representations, computational reformulations, and out-of-sample performance guarantees.

DRO is not new to insurance. Recent work has explicitly considered distributionally robust insurance using Wasserstein ambiguity sets. Boonen and Jiang developed a framework in which ambiguity around a benchmark loss distribution is represented by a Wasserstein ball and analyzed optimal insurance contracting under distributional uncertainty. More recently, Sliwinski, Llamazares-Elias, Siska, and Szpruch investigated phi-divergence-based DRO for offline insurance pricing and reported that robust policies can become overly conservative while producing limited performance improvements under distributional shifts.

These developments suggest that the important research question is no longer whether DRO can be applied to insurance. Instead, a more focused question concerns how the uncertainty produced by modern probabilistic risk models should determine the ambiguity set used for robust pricing.

This paper develops that research perspective.

2. Background and Related Literature

2.1. Machine Learning in Insurance Pricing

Machine learning has become an important component of modern actuarial modeling. Its principal attraction is the ability to approximate complex relationships without imposing the relatively restrictive functional structures associated with many traditional actuarial models. Tree-based methods, boosting algorithms, neural networks, and other flexible learning procedures can capture nonlinearities and interactions among risk factors. [3]

The literature nevertheless demonstrates that predictive performance alone is insufficient for insurance pricing. Insurance premiums are not merely predictions of an outcome; they are financial decisions that must satisfy additional requirements concerning balance, calibration, interpretability, stability, and economic adequacy.

Blier-Wong et al. reviewed the development of machine learning applications in property and casualty insurance and showed that research has increasingly examined flexible algorithms for ratemaking and reserving.This development provides an important foundation for the present research, but it also highlights the need to distinguish between flexible prediction and reliable decision-making.

Denuit et al. focused specifically on the calibration of machine-learning insurance pricing models. Their work on auto-calibration demonstrates that models trained for predictive performance may require additional calibration procedures to achieve desirable balance properties.

The present research builds on this literature but shifts the question from calibration of the premium prediction alone to calibration of predictive uncertainty for robust decision-making.

2.2. Distributionally Robust Optimization

DRO addresses optimization problems in which the probability distribution governing uncertain quantities is not known exactly. Instead of optimizing against a single probability model, the decision-maker considers an ambiguity set of distributions consistent with available statistical or structural information. [4]

The development of Wasserstein-based DRO has been particularly influential because Wasserstein distances provide a meaningful way of measuring discrepancies between probability distributions. Esfahani and Kuhn established important performance guarantees and tractable reformulations for data-driven DRO problems based on Wasserstein ambiguity sets. Their work helped establish a rigorous connection between finite-sample statistical information and robust optimization.

A recent broad survey by Kuhn, Shafiee, and Wiesemann emphasizes the central role of ambiguity sets in modern DRO and its connections with statistics, optimization, control, and machine learning.

For insurance research, these developments are particularly relevant because the loss distribution is rarely known with certainty. Historical claims data provide information about the distribution, but future losses may differ from historical observations because of sampling variability, changing environments, model misspecification, or other sources of uncertainty.

2.3. Distributionally Robust Insurance

Recent research has established that distributional ambiguity can be incorporated directly into insurance problems.

Boonen and Jiang considered optimal insurance under ambiguity represented by Wasserstein balls around a benchmark loss distribution. Their analysis provides mathematical characterizations of worst-case distributions and optimal insurance structures.

Sliwinski et al. considered phi-divergence-based DRO for offline insurance pricing and reported an important practical finding: robust policies may become overly conservative, while their gains under distributional shifts may remain limited.

This observation is particularly relevant to the present research because it identifies a fundamental tension. Robustness is valuable when the predictive distribution is uncertain, but robustness obtained by expanding the ambiguity set indefinitely is not necessarily desirable. A large ambiguity set can produce premiums that are unnecessarily high or otherwise economically unattractive.

The research problem therefore concerns not merely robust pricing, but calibrated robust pricing. [5]

3. Research Gap and Objectives

The literature reviewed above indicates that machine-learning-based insurance pricing, predictive calibration, distributionally robust optimization, and distributionally robust insurance are each established research areas.

The more specific gap lies at their intersection.

A modern AI-based insurance model may provide a conditional predictive distribution rather than a single estimate. That predictive distribution contains information about uncertainty. However, the uncertainty information produced by the prediction system does not automatically determine the ambiguity set required by a DRO problem.

An ambiguity set has at least two important characteristics: its center or reference distribution and its size or geometry. The reference distribution may be generated by an AI-based probabilistic model. The more difficult question concerns the size and structure of the neighborhood surrounding that distribution.

If the ambiguity set is too narrow, the robust premium may fail to reflect important model uncertainty. If it is too broad, the robust premium may become excessively conservative. The resulting problem is therefore one of calibration.

The central research gap can be expressed as follows:

How should predictive uncertainty generated by an AI-based insurance risk model be statistically translated into the size and structure of a distributional ambiguity set, and how does this translation determine the robustness and conservatism of the resulting premium?

This formulation is deliberately narrower than a general claim that AI and DRO have not previously been combined. Existing literature demonstrates that both areas have already been connected in various forms. The contribution proposed here is instead to investigate the mathematical calibration mechanism connecting predictive uncertainty to distributional ambiguity in insurance pricing.

Research Objectives

Objective 1: Probabilistic AI Risk Modeling and Uncertainty Quantification

The first objective is to investigate probabilistic machine-learning models capable of representing conditional insurance loss distributions. [6]

The focus will not be on creating a new generic neural network architecture. Instead, AI will be treated as a flexible mechanism for estimating the parameters or structure of a probabilistic loss model.

The research will examine how predictive uncertainty can be represented in a manner suitable for mathematical analysis. Candidate approaches may include distributional regression, ensemble-based uncertainty estimates, or Bayesian approximations.

An important requirement is that uncertainty measures should be interpretable in probabilistic terms and should be capable of being connected to subsequent statistical guarantees.

The first objective will therefore establish the mathematical interface between flexible AI prediction and probabilistic insurance risk modeling.

Objective 2: Statistically Calibrated Ambiguity Sets for Robust Pricing

The second objective is the central methodological component.

Given a predictive distribution and a measure of predictive uncertainty, the research will investigate how to construct an ambiguity set whose size and geometry are statistically justified.

Wasserstein-type sets provide one possible framework because they allow the distance between the nominal predictive distribution and alternative distributions to be quantified. Divergence-based alternatives may also be considered where mathematically appropriate.

The central issue is not simply the selection of an ambiguity-set family. Instead, the research will ask whether its parameters can be related systematically to predictive uncertainty, calibration quality, and available data.

The desired outcome is an ambiguity set that is neither artificially narrow nor unnecessarily broad.

The resulting robust premium should therefore reflect a controlled degree of distributional uncertainty.

Objective 3: Robustness, Stability, and Conservatism

The third objective is to characterize the consequences of ambiguity-set construction for insurance premiums. [7]

Particular attention will be given to three properties.

First is robustness: how well does the premium remain protected when the true loss distribution differs from the nominal predictive distribution?

Second is stability: how sensitive is the robust premium to small changes in the predictive distribution or uncertainty estimate?

Third is conservatism: how much does the robust premium differ from the nominal premium as the ambiguity set expands?

These questions can potentially lead to theoretical results concerning continuity, sensitivity, convergence, and upper or lower bounds for robust premiums.

The project will therefore seek a mathematical characterization of the relationship:

predictive uncertainty → ambiguity size → robustness → conservatism.

4. Proposed Conceptual Framework

The proposed framework consists of four logically connected stages. [8]

The first stage is probabilistic risk prediction. An AI-based model estimates a conditional distribution of insurance losses based on observed risk characteristics.

The second stage is uncertainty quantification. Instead of treating the estimated predictive distribution as known, the framework characterizes uncertainty surrounding the prediction. Such uncertainty may arise from finite samples, parameter estimation, model specification, or other sources.

The third stage converts predictive uncertainty into a statistically calibrated ambiguity set. The ambiguity set represents distributions considered plausible given the information contained in the predictive model and its uncertainty assessment.

The fourth stage determines a robust insurance premium using the ambiguity set. The premium is therefore designed not only for the nominal predicted distribution but also for a specified neighborhood of plausible distributions.

The conceptual chain is:

AI-based probabilistic prediction → predictive uncertainty → calibrated ambiguity → robust premium.

The central research problem is the mathematical characterization of the arrows connecting these stages.

5. Theoretical Research Direction

The mathematical component of the research can be organized around several classes of results.

The first concerns the consistency of uncertainty estimates and predictive distributions. Under appropriate assumptions, the research may investigate whether estimated predictive distributions converge toward a limiting distribution as the amount of available information increases. [9]

The second concerns the convergence or stability of uncertainty-derived ambiguity sets. If predictive uncertainty becomes smaller as information increases, an important question is whether the associated ambiguity sets contract in a mathematically controlled manner.

The third concerns robust pricing functionals. The research may investigate conditions under which robust premiums are continuous or Lipschitz with respect to changes in the reference predictive distribution or ambiguity parameter.

The fourth concerns conservatism bounds. A central objective is to determine whether the difference between a robust and nominal premium can be bounded in terms of a statistically meaningful measure of uncertainty.

The fifth concerns calibration. If predictive calibration improves, the framework will investigate whether smaller ambiguity sets can provide comparable robustness guarantees and thereby reduce unnecessary conservatism.

These questions are mathematically meaningful because they involve probability distributions, statistical estimation, optimization, risk functionals, and sensitivity analysis.

6. Scientific and Practical Significance

A fundamental motivation for this research is the tension between robustness and conservatism. [10]

A non-robust pricing model treats the estimated predictive distribution as sufficiently reliable. This may expose the insurer to model risk if the estimated distribution differs materially from the future loss-generating process.

A fully conservative approach could instead consider a very large family of distributions. Such a procedure might protect against many forms of distributional misspecification but could produce premiums that are excessively high.

The problem is therefore not to maximize robustness without qualification.

The objective is to determine an appropriate level of robustness based on measurable uncertainty.

This perspective is supported by recent insurance DRO research. Sliwinski et al. found that phi-divergence-based robust insurance pricing may result in overly conservative policies and limited gains under distributional shifts. This result makes the study of conservatism more than a theoretical preference; it is a practical methodological issue.

The proposed framework therefore treats conservatism as an explicit research quantity rather than an undesirable side effect that is considered only after model development.

The proposed research does not introduce machine learning to insurance pricing as a new idea. Nor does it introduce DRO to insurance.

Its potential contribution is a more specific mathematical framework for connecting these existing components. [11]

The principal scientific contribution would be the formalization of a mechanism through which predictive uncertainty can determine distributional ambiguity for robust pricing.

Such a framework could contribute to actuarial mathematics by providing a more systematic treatment of model uncertainty in premium determination.

It could contribute to mathematical statistics by connecting predictive calibration with distributional neighborhoods.

It could contribute to optimization by investigating robust pricing functionals whose ambiguity sets are derived from predictive rather than purely empirical uncertainty.

Finally, it could contribute to AI-based risk modeling by demonstrating that predictive uncertainty can have a direct mathematical role in downstream decision-making.

The practical importance of the framework is associated with the increasing complexity of insurance data and the growing use of algorithmic risk models.

A model that provides a highly accurate prediction but fails to quantify uncertainty may provide insufficient information for robust decision-making. Conversely, a robust optimization method that ignores the quality of the underlying predictive model may introduce unnecessary conservatism.

The proposed framework attempts to connect these two perspectives.

For an insurer, the eventual objective would be a pricing methodology in which the robustness adjustment is related to measurable predictive uncertainty rather than being selected solely through arbitrary stress margins.

Such an approach could potentially make robust pricing more transparent because the additional conservatism in a premium would have a mathematical interpretation.

However, practical adoption would require further investigation of regulatory, fairness, interpretability, computational, and market considerations. The present research is therefore primarily methodological and mathematical rather than a claim that the proposed framework is immediately deployable in commercial insurance systems. [12]

7. Limitations and Future Research

The project should remain deliberately focused.

It does not attempt to develop a universal AI system for all insurance products. It does not attempt to address simultaneously pricing, underwriting, reserving, capital allocation, systemic risk, climate risk, and reinsurance.

A focused non-life insurance pricing setting provides a more appropriate environment for developing and analyzing the mathematical framework.

Likewise, the research should not assume that the proposed approach will outperform every existing actuarial or machine-learning method. [13]

The central objective is to establish theoretical relationships and then evaluate their implications through controlled numerical experiments and, where suitable data are available, empirical illustrations.

This distinction is important because recent research has shown that robust methods can sometimes offer limited gains under distributional shifts despite their theoretical robustness.

Several extensions could follow from the proposed framework.

One direction would be the development of conditional ambiguity sets whose size depends on individual risk characteristics rather than being identical across the entire portfolio.

Another would involve dynamic insurance pricing, where predictive uncertainty evolves over time as new claims and policy information become available.

A further extension would consider fairness and regulatory constraints. These issues are increasingly important in AI-based insurance pricing, but incorporating them into the present mathematical framework would substantially broaden the project.

Finally, the methodology could be extended to other risk-sensitive domains such as finance, energy, or healthcare. Such extensions should, however, be considered only after the core insurance pricing problem has been mathematically established.

8. Conclusion

Machine learning has expanded the capacity of insurance pricing models to capture complex relationships in high-dimensional data. At the same time, probabilistic AI models introduce an opportunity to represent predictive uncertainty more explicitly. Distributionally robust optimization provides a mathematically rigorous framework for making decisions when the probability distribution itself is uncertain. [14]

The combination of these developments creates a promising research direction, but it also creates a central methodological challenge: predictive uncertainty must somehow be translated into a mathematically meaningful ambiguity set, and the resulting ambiguity must be sufficiently robust without producing unnecessary conservatism.

Existing research has already established the foundations of machine-learning-based insurance pricing, predictive calibration, Wasserstein-based distributionally robust optimization, and robust insurance. The proposed research therefore does not claim novelty for any of these components independently.

Instead, it focuses on their intersection: the statistical and mathematical calibration of distributional ambiguity from AI-derived predictive uncertainty and the resulting relationship between robustness and premium conservatism.

This perspective offers a focused research program at the intersection of actuarial mathematics, mathematical statistics, artificial intelligence, and optimization. Its central challenge is not merely to construct a robust premium, but to understand mathematically how much robustness is justified by the available predictive uncertainty and how much conservatism that robustness necessarily introduces. [15]

A successful theoretical framework would provide a foundation for uncertainty-aware insurance pricing in which the robustness adjustment is statistically interpretable, mathematically analyzable, and explicitly connected to the quality of the underlying predictive model.

Abbreviations

AI	Artificial Intelligence
DRO	Distributionally Robust Optimization
P&C	Property and Casualty

Author Contributions

Majid Ghorbani: Conceptualization, Formal Analysis, Investigation, Metodology, Data curation, Supervision, Writing – original draft, Writing – review & editing

Conflicts of Interest

The author declares no conflicts of interest.

References

[1] Blier-Wong, C., Cossette, H., Lamontagne, L., & Marceau, E. (2021). Machine Learning in P&C Insurance: A Review for Pricing and Reserving. Risks, 9(1), 4. https://doi.org/10.3390/risks9010004

[2] Denuit, M., Charpentier, A., & Trufin, J. (2021). Autocalibration and Tweedie-dominance for insurance pricing with machine learning. Insurance: Mathematics and Economics, 101, 485–497. https://doi.org/10.1016/j.insmatheco.2021.09.001

[3] Mohajerin Esfahani, P., & Kuhn, D. (2018). Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. Mathematical Programming, 171, 115–166. https://doi.org/10.1007/s10107-017-1172-1

[4] Kuhn, D., Shafiee, S., & Wiesemann, W. (2025). Distributionally robust optimization. Acta Numerica, 34, 579–804. https://doi.org/10.1017/S0962492924000084

[5] Ghorbani, M. (2025). Geometric Design of Bread in Iran: A Mathematical Approach to Waste Reduction and Supply Chain Efficiency. Mathematical Modelling and Applications, 10(3), 43-48. https://doi.org/10.11648/j.mma.20251003.11

[6] Boonen, T. J., & Jiang, W. (2025). Distributionally robust insurance under the Wasserstein distance. Insurance: Mathematics and Economics, 120, 61–78. https://doi.org/10.1016/j.insmatheco.2024.11.003

[7] Sliwinski, L., Llamazares-Elias, L., Siska, D., & Szpruch, L. (2025). Parametric Phi-Divergence-Based Distributionally Robust Optimization for Insurance Pricing. Proceedings of the 6th ACM International Conference on AI in Finance, 378–386. https://doi.org/10.1145/3768292.3770404

[8] Zhang, Y., Ji, L., Aivaliotis, G., & Taylor, C. C. (2024). Bayesian CART models for insurance claims frequency. Insurance: Mathematics and Economics, 114, 108–131. https://doi.org/10.1016/j.insmatheco.2023.11.005

[9] Finger, D., Albrecher, H., & Wilhelmy, L. (2024). On the cost of risk misspecification in insurance pricing. Japanese Journal of Statistics and Data Science, 7, 1111–1153. https://doi.org/10.1007/s42081-024-00256-z

[10] Denuit, M., Trufin, J., & others. (2024). Convex and Lorenz orders under balance correction in nonlife insurance pricing: Review and new developments. Insurance: Mathematics and Economics, 118, 123–128. https://doi.org/10.1016/j.insmatheco.2024.06.003

[11] Lindholm, M., & Palmquist, J. (2024). Black-box guided generalised linear model building with non-life pricing applications. Annals of Actuarial Science, 18(3), 675–691. https://doi.org/10.1017/S1748499524000265

[12] Lindholm, M., & Nazar, T. (2024). On duration effects in non-life insurance pricing. European Actuarial Journal, 14, 809–832. https://doi.org/10.1007/s13385-024-00385-5

[13] Dong, P., & Quan, Z. (2025). Automated machine learning in insurance. Insurance: Mathematics and Economics, 120, 17–41. https://doi.org/10.1016/j.insmatheco.2024.10.002

[14] A Neural Network Approach for Pricing Correlated Health Risks. (2025). Risks, 13(5), 82. https://doi.org/10.3390/risks13050082

[15] Lindholm, M., Richman, R., Tsanakas, A., & Wüthrich, M. V. (2024). What is fair? Proxy discrimination vs. demographic disparities in insurance pricing. Scandinavian Actuarial Journal, 935–970. https://doi.org/10.1080/03461238.2024.2364741