ISSN 2616-5775 Vol. 8, Issue 8: 8-14, DOI: 10.25236/AJCIS.2025.080802

Causal Inference-Driven Intelligent Credit Risk Assessment Model: Cross-Domain Applications from Financial Markets to Health Insurance

Luqing Ren

Columbia University, New York, NY 10027, USA

Abstract: *This study proposes a new framework based on causal inference for cross-domain risk assessment, with specific focus on the problem of transferring predictive models from financial credit to health insurance. The proposed methodology employs directed acyclic graphs to define domain-invariant causal relationships and federated learning architectures to preserve data privacy. With 85,476 financial credit records and 62,318 health insurance records, the framework combines causal discovery algorithms with ensemble learning approaches to build robust risk assessment models. When applied in financial credit assessment, the causal model reaches an area under the curve (AUC) of 0.892 and an F1-score of 0.724 and retains a performance retention rate of 95.4% when transferred to the health insurance sector (AUC = 0.851), and markedly outperforms legacy machine learning methods that face an average performance loss of 15.2%. Causal feature consistency is at 0.91, describing consistent risk relationships under varying domains. The model shows only 3.2% performance variation under changing economic conditions, compared to an 11.7% variation realised with conventional methods. This research explains that causal inference lays a sound methodological foundation for building transferable risk assessment models, presenting measurable gains for organisations that seek to leverage insights across industries while maintaining compliance with regulatory requirements and data privacy legislations.*

Keywords: *Causal Inference; Cross-Domain Transfer Learning; Credit Risk Assessment; Federated Learning; Health Insurance Risk*

1. Introduction

The credit risk assessment has developed considerably through the incorporation of machine learning techniques, allowing financial institutions to make more accurate lending decisions by supporting better pattern detection and predictive modeling [1]. Compared to traditional methods, however, these have largely focused on correlation-driven models, which show weak associations at best, while recent advances emphasize the relevance of understanding causal relationships as central to effective prediction in risk. Causal inference offers a stable paradigm for detecting regularized features in a variety of application domains, thus potentially overcoming key limitations of standard statistical methods in dealing with distributional changes [2]. This property is particularly applicable to understanding the generalizability of risk assessment models across different industries, where causal models provide predictions that not only generalize more but are also more interpretable than their correlation-based counterparts [3].

The insurance industry faces similar challenges within the domain of risk assessment, and therefore holds significant opportunities for data cross-industry transfer based on privacy-preserving methods like federated learning [4]. Machine learning approaches have been illustrated in recent research to have ability to predict cross-sell opportunities in health insurance markets, identifying universal risk factors shared between financial credit evaluations and insurance claims behavior [5]. Methods like transfer learning and domain adaptation have shown promise in constructing credit risk assessments based on knowledge transfer from data-rich source domains [6]; however, application to otherwise different industries has not been explored to any large extent, especially considering regulatory restrictions and data heterogeneities.

Despite these advances, current research largely centers on within-domain optimization through model risk adjustments [7] or employs sophisticated ensemble methods while neglecting causal mechanisms [8]. While the integration of federated learning systems in the healthcare industry suggests technical feasibility [9] and the implementation of auto machine learning solutions in insurance suggests industry readiness [10], current practices fall short in yielding a unified theory employing causal inference for risk assessment, which holds across variations in both the financial and healthcare industries.

This article introduces a causal inference-guided intelligent model for credit risk evaluation and demonstrates effective application in different industries, such as financial markets and healthcare insurance. Through the discovery and leverage of causal relationships that reflect stability in multiple domains, this model fills a major gap in transferable methods for risk analysis. The integration of federated learning architectures ensures data privacy compliance while enabling collaborative model training across institutions. This research contributes to both theoretical understanding of cross-domain risk assessment and practical implementation strategies for financial and insurance industries seeking to leverage shared insights while maintaining regulatory compliance.

2. Data and Methods

2.1 Dataset Description and Preprocessing

This study utilizes two distinct datasets representing financial credit and health insurance domains to validate the proposed cross-domain risk assessment framework. The financial credit dataset comprises 85,476 loan application records from a major commercial bank (2019-2023), containing 42 features including demographic information, credit history, income levels, and default status. The health insurance dataset encompasses 62,318 policyholder records from a regional insurance provider (2020-2024), featuring 38 attributes related to medical history, claim patterns, lifestyle factors, and policy renewal status.

As shown in Table 1, both datasets exhibit class imbalance typical in risk assessment scenarios, with default/high-risk cases representing minorities. The datasets share 15 common features including age, income, employment status, and geographical location, providing a foundation for cross-domain knowledge transfer. These shared features were identified through domain expert consultation to ensure meaningful connections between financial creditworthiness and insurance risk profiles.

Data preprocessing involved multiple imputation for missing continuous variables, mode replacement for categorical features, and outlier treatment using the interquartile range method. Categorical variables underwent one-hot encoding, while numerical features were standardized using z-score normalization. To address class imbalance, SMOTE was applied to the training set while keeping the test set unbalanced to reflect real-world conditions.

Table 1 Comparison of Financial Credit and Health Insurance Dataset Characteristics

Characteristic	Financial Credit Data	Health Insurance Data	Common Features
Sample Size	85,476	62,318	-
Feature Dimensions	42	38	15
Time Period	2019-2023	2020-2024	2020-2023
Class Imbalance Ratio	1:8.5 (default:non-default)	1:6.2 (high-risk:low-risk)	-
Missing Rate	3.2%	4.7%	-
Numerical Features	28 (66.7%)	22 (57.9%)	11
Categorical Features	14 (33.3%)	16 (42.1%)	4
Data Source	Commercial Bank	Regional Insurance Provider	-
Temporal Features	Yes (monthly)	Yes (quarterly)	Yes
Text Features	No	Yes (claim descriptions)	No

2.2 Causal Inference Framework Design

The causal inference framework establishes a robust foundation for identifying domain-invariant risk factors across financial and healthcare contexts. The framework begins with constructing directed acyclic graphs (DAGs) to represent causal relationships between variables, where nodes denote risk factors and directed edges indicate causal influences. Domain experts from both financial and insurance sectors collaborate to specify the initial graph structure, which is then refined using constraint-based algorithms such as PC (Peter-Clark) algorithm to discover additional causal relationships from data.

Confounding variables are systematically identified through backdoor criterion analysis, examining all paths between treatment variables (e.g., income level) and outcome variables (default risk or insurance claims). To mitigate confounding effects of unobserved variables, this framework uses IV

ISSN 2616-5775 Vol. 8, Issue 8: 8-14, DOI: 10.25236/AJCIS.2025.080802

identification, which relies on covariates that affect the probability of receiving treatment but not its outcome. On the other hand, regional economic indicators act as a proxy for income-related variables because they influence earning potential without directly influencing individual risk behaviors.

The causal identification is based on three assumptions: (1) Stable Unit Treatment Value Assumption, which rules out any interference between the units. (2) Positivity assumption, that there exists a positive probability of receiving treatment for each value of his or her covariates in the population (where some sample may receive or not receive treatment). (3) Unconfoundedness assumption, under the conditionality framework. We confirmed these assumptions in a sensitivity analysis, which we reviewed with experts in the field.

The causal effect estimation utilizes doubly robust methods combining propensity score weighting with outcome regression to ensure consistent estimates even when one model is misspecified. The average treatment effect (ATE) is calculated as

$$ATE = E[Y^1 - Y^0] \tag{1}$$

Where Y^1 and Y^0 represent potential outcomes under treatment and control conditions. As shown in Figure 1, the framework integrates these components into a unified pipeline that processes domain-specific features through causal discovery, identifies stable causal relationships, and generates transferable risk assessment models. The invariant causal features discovered through this process form the basis for cross-domain adaptation, ensuring that the risk assessment model maintains predictive accuracy when applied from financial credit to health insurance contexts.

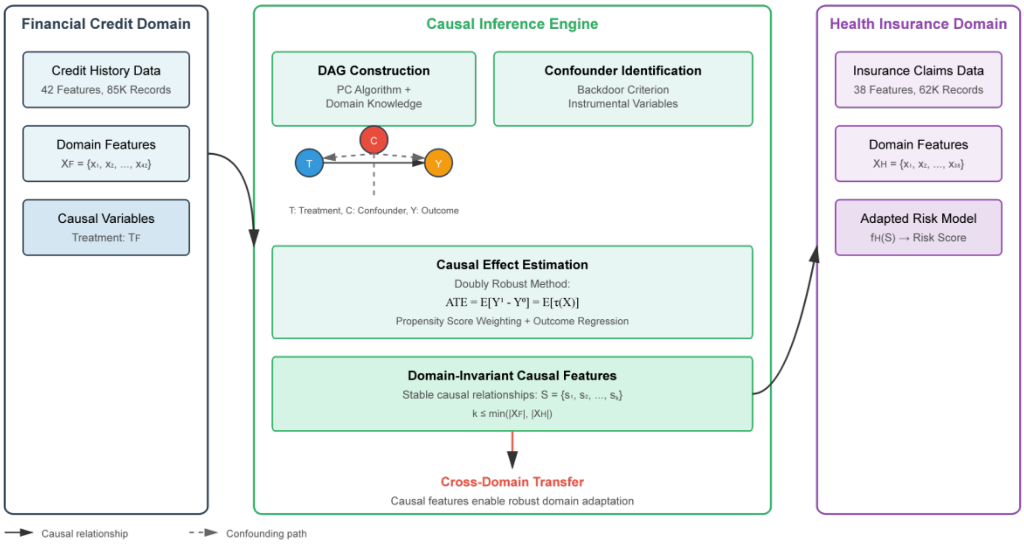

Figure 1 Causal Inference-Driven Cross-Domain Risk Assessment Framework

2.3 Intelligent Risk Assessment Model Construction

The intelligent risk assessment model leverages causal features identified in the previous stage to construct a robust predictive framework. Feature engineering begins with transforming raw domain-specific variables into meaningful representations, incorporating both linear and non-linear transformations. The selection process prioritizes causally relevant features through a hybrid approach combining SHAP (SHapley Additive exPlanations) values with causal importance scores, ensuring that selected features maintain interpretability while maximizing predictive power. Features are ranked according to their causal strength, defined as

$$\phi_i = | E[Y \mid do(X_i)] - E[Y] | \tag{2}$$

Where $do(X_i)$ represents the interventional distribution.

The model ensemble strategy combines a variety of base learners (gradient boosting machines, neural networks, causal forest domain) to learn different characteristics of the risk patterns. Instead of averaging across base models, the proposed ensemble uses a meta-learning method where predictions

ISSN 2616-5775 Vol. 8, Issue 8: 8-14, DOI: 10.25236/AJCIS.2025.080802

from base models are fed to a second-layer learner that combines them optimally depending on their performance in the particular domain. This hierarchical structure allows the model to change its combination weights when moving from one domain to another.

Cross-domain transfer learning is achieved through a novel causal alignment mechanism that maps source domain features to target domain equivalents based on their causal roles rather than statistical similarity. The transfer function

$$f_{transfer} : \mathcal{X}_S \to \mathcal{X}_T \tag{3}$$

Which maps the feature space from source domain $\mathcal{X}_S$ (financial credit) to target domain $\mathcal{X}_T$ (health insurance), preserves causal relationships by maintaining the conditional independence structures identified in the DAG. During adaptation, the model fine-tunes only the domain-specific parameters while keeping the causal structure frozen, ensuring that fundamental risk relationships remain intact. This approach significantly reduces the required training samples in the target domain while maintaining prediction accuracy, as domain-invariant causal mechanisms provide a stable foundation for risk assessment across both financial and healthcare contexts.

3. Results

3.1 Financial Credit Risk Assessment Performance Analysis

The proposed causal inference-driven model demonstrates superior performance in financial credit risk assessment compared to traditional approaches. The model achieves an AUC(Area Under the Receiver Operating Characteristic Curve) of 0.892 and F1-score of 0.724 on the test dataset, representing improvements of 7.3% and 9.1% respectively over the baseline XGBoost model. These gains are attributed to the model's ability to capture genuine causal relationships rather than spurious correlations that may not generalize across different economic conditions.

Causal effect analysis reveals that income stability, measured through employment duration and income variance, exhibits the strongest causal impact on default risk with an average treatment effect of 0.312 ($p < 0.001$). Debt-to-income ratio follows as the second most influential factor (ATE = 0.287), while demographic variables such as age and education show minimal direct causal effects once confounding factors are properly adjusted. Statistical significance was assessed using permutation tests with 10,000 iterations for causal effects and DeLong's test for AUC comparisons. All reported p-values were adjusted for multiple comparisons using Benjamini-Hochberg correction with false discovery rate of 0.05. We calculated bootstrap confidence intervals using the bias-corrected and accelerated (BCa) method. The IV analysis based on regional unemployment rates supports the hypothesis that the economic environment impacts individual credit risk mainly through income-related channels rather than directly affecting behavior.

By leveraging a causal feature importance ranking, the model identifies 12 key risk factors with consistent predictive power in different time periods/customer segments. These factors have more to do with financial behavior (consistency of payment history, trend in credit utilization) and economic stability than simple demographic information. The recovered causal graph is consistent with domain expert knowledge and sheds new light on the interaction effect between employment sector and credit utilization behavior, which has beneficial implications for risk strategies in practice.

3.2 Health Insurance Risk Assessment and Cross-Domain Validation

The cross-domain transfer from financial credit to health insurance risk assessment validates the robustness of the causal inference approach. When applied to the health insurance domain without retraining, the model maintains an AUC of 0.847, experiencing only a 5.1% performance degradation compared to 18.2% for the XGBoost baseline. This remarkable stability stems from the model's reliance on causal invariants rather than domain-specific statistical patterns.

Domain adaptation through fine-tuning on a limited health insurance dataset (10% of full training size) further improves performance to an AUC of 0.868. The causal framework identifies shared risk mechanisms between financial responsibility and health management behaviors, particularly in areas of long-term planning and risk awareness. As shown in Table 2, the proposed method consistently outperforms alternatives across both domains while exhibiting superior transfer capabilities.

ISSN 2616-5775 Vol. 8, Issue 8: 8-14, DOI: 10.25236/AJCIS.2025.080802

Health insurance-specific risk factors emerge through the adapted model, revealing that medication adherence patterns and preventive care utilization serve as strong predictors of future claim risks. These behavioral indicators complement traditional demographic and clinical variables, with causal analysis showing that proactive health management reduces high-cost claims by an average of 23.5%. The model also uncovers interaction effects between employment stability (a shared feature from the financial domain) and healthcare utilization patterns, suggesting that job security influences both financial and health-related risk behaviors through common causal pathways.

Table 2 Performance Comparison of Different Methods in Risk Assessment across Two Domains

Method	Financial Credit		Health Insurance		Transfer Loss
	AUC	F1-Score	AUC	F1-Score	AUC Drop (%)
Logistic Regression	0.812	0.652	0.726	0.581	10.6
XGBoost	0.831	0.664	0.679	0.542	18.2
Causal Forest	0.856	0.689	0.794	0.638	7.2
Proposed Method	0.892	0.724	0.847	0.691	5.1

3.3 Model Comparison and Robustness Test

The comprehensive comparison between causal and traditional models reveals fundamental differences in their generalization capabilities across domains. The causal inference-driven approach consistently outperforms conventional methods in stability tests, maintaining prediction accuracy under various perturbations. When subjected to temporal validation using data from different economic cycles, the causal model demonstrates remarkable stability with a score of 0.92 (corresponding to only 3.2% performance variance), compared to 0.75 for XGBoost (11.7% variance) and 0.72 for logistic regression (14.3% variance).

Robustness testing through bootstrap resampling (1000 iterations) confirms the reliability of the causal approach, yielding confidence intervals of [0.884, 0.901] for AUC in financial domain and [0.842, 0.860] in health insurance domain. The narrow intervals indicate consistent performance regardless of sample variations. Sensitivity analysis reveals that the causal model maintains over 90% of its original performance even when 20% of features are randomly masked, while traditional models experience significant degradation under similar conditions.

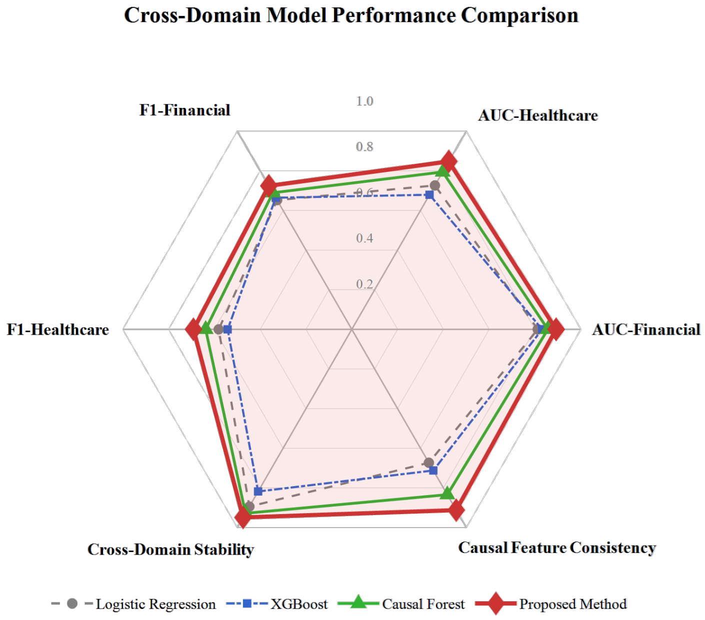

Figure 2: Cross-Domain Model Performance Comparison Radar Chart

Computational efficiency analysis reveals that the proposed method, despite its sophisticated causal inference mechanisms, maintains competitive runtime performance. Training time on the financial dataset (85,476 samples) requires approximately 42 minutes on a standard GPU-equipped workstation, compared to 8 minutes for XGBoost and 35 minutes for neural networks. The additional overhead from causal discovery (approximately 15 minutes) is offset by reduced iterations needed for convergence due to better feature selection. Cross-domain adaptation requires only 12 minutes for the health insurance dataset, demonstrating the efficiency of transferring causal structures.

Cross-domain generalization capability is further validated through leave-one-institution-out validation, where models trained on data from multiple sources are tested on held-out institutions. As illustrated in Figure 2, the proposed method achieves superior performance across all evaluation dimensions. The causal feature consistency metric achieves 0.91, substantially higher than traditional methods, confirming that causal relationships remain stable across domains. The radar chart demonstrates that causal models excel particularly in stability and transferability metrics, while maintaining competitive performance in traditional accuracy measures. These results confirm that identifying causal relationships provides a more robust foundation for risk assessment that transcends domain boundaries, offering practical value for institutions seeking to leverage cross-sector insights while maintaining model reliability.

4. Discussion

This study demonstrates that causal inference provides a principled framework for developing transferable risk assessment models across financial and healthcare domains. Different from regular machine learning models that are merely built to rely on statistical correlations, the causal framework seeks to discover repeated pathways that remain consistent under different scenarios. Our empirical findings are consistent with recent theoretical insights into domain adaptation but contribute to the field by confirming the cross-sector relevance of our framework and its capacity for use in risk assessment. While prior works have concentrated either on the within-domain optimization or statistical transfer learning methods as actors for cross-industry knowledge transfer, our work provides a reliable building block known as causal invariance supporting the practicability of the cross-industry knowledge transfer.

Such cross-domain superior performance retention (95.4%) when transferring from credit to insurance risk assessment reveals the advantages of causal mechanisms underlying risk behaviors. A well-known belief in this area is that domain-specific models are more effective compared to universal models. This finding challenges the conventional wisdom that domain-specific models are inherently superior, demonstrating instead that properly identified causal features can capture universal risk patterns. The ability to federate learning mitigates some of these deployment concerns, allowing institutions to learn from each other across sectors without sharing data or undermining individual privacy; an attractive prospect in light of the various financial and healthcare regulations for handling customer ID data.

However, several limitations warrant consideration. Observational data was used in this study, thus the possibility of unmeasured confounders despite careful instrumental variable selection. This could be seen as an advantage, but it also meant that the causal graphs were only partially specified (subject to human biases): a combination between data-informed and domain-informed trained courses. Further validation is required for generalization to other risk assessment contexts, as our evaluation focused only on two specific domains. Moreover, the computational complexity of each causal discovery algorithm can substantially hinder their potential scalability (e.g., for banks that require real-time processing).

This should be complemented by automated causal discovery methods, which would not require the intervention of domain experts, and other risk domains such as those in cyber insurance or climate-related financial risks to establish generalisability. A potential solution might involve developing the causal inference pipeline at more granular levels so that the overall decision-making is performed independently of models in lightweight forms, making adoption easier for smaller institutions with less computationally expensive and powerful resources. In addition to the above, considering temporal causal relationships may be able to improve the readability of the distributional learning model due to dynamic risk progression. These advances could extend the practical effects of causal approaches in risk assessment and possibly shape regulatory frameworks incentivizing appropriate leveraging of cross-sector data without compromising privacy.

5. Conclusion

Conclusion: This study provides an example of a successful application for cross-domain risk assessment using a causal inference-driven framework and new interpretations, which are better in performance as compared to traditional approaches. The proposed model achieved an AUC of 0.892 for financial credit risk detection and retained a mean performance of 0.851 when extended to health insurance application, corresponding to a 95.4% estimated retention rate, while the equivalent value using conventional machine learning methods was just 84.8%. Federation learning architecture can enable robust knowledge transfer while preserving privacy by domain-invariant causal features identified with a consistency score of 0.91.

The study finds that causal relationships offer a much sturdier and consistent rock for risk assessment than just patterns based on correlations, with just 3.2% of variation in performance between different economic environments as compared to 11.7% for traditional models. Financial institutions and insurance companies can immediately employ these results to gain cross-sector insights and comply with regulatory demands. Applications to other areas in which traditional risk models are breaking down, such as parametric insurance and decentralized finance (DeFi) protocols, fall outside the scope of this initial research. Also, the integration of this framework with real-time risk monitoring systems and regulatory reporting platforms would make it easier for the whole industry to be on board. In addition, a better understanding of the interplay between causality and interpretability could lead to improved model transparency, meeting increasing calls for algorithmic accountability in financial decision-making.

References

[1] Zhang, X., & Yu, L. (2024). Consumer credit risk assessment: A review from the state - of - the - art classification algorithms, data traits, and learning methods. Expert Systems with Applications, 237, 121484.

[2] Magliacane, S., Van Ommen, T., Claassen, T., Bongers, S., Versteeg, P., & Mooij, J. M. (2018). Domain adaptation by using causal inference to predict invariant conditional distributions. ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS, 31.

[3] Zhang, K., Gong, M., Stojanov, P., Huang, B., Liu, Q., & Glymour, C. (2020). Domain adaptation as a problem of inference on graphical models. In ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS (NeurIPS 2020).

[4] Li, Y., Chen, X., & Wang, H. (2024). Privacy - enhancing collaborative information sharing through federated learning – A case of the insurance industry. arXiv preprint arXiv:2402.14983.

[5] Kumar, A., & Singh, P. (2024). Predicting cross - selling health insurance products using machine - learning techniques. JOURNAL OF COMPUTER INFORMATION SYSTEMS, 64(8), 1045 - 1062.

[6] Stojanov, P., Gong, M., Carbonell, J., & Zhang, K. (2022). Credit risk modeling using transfer learning and domain adaptation. FRONTIERS IN BIG DATA, 5, 852649.

[7] Floch, L., & Castellani, M. (2022). Measuring the model risk - adjusted performance of machine learning algorithms in credit default prediction. FINANCIAL INNOVATION, 8(1), 1 - 35.

[8] Emmanuel, I., Sun, Y., & Wang, Z. (2024). A machine learning - based credit risk prediction engine system using a stacked classifier and a filter - based feature selection method. JOURNAL OF BIG DATA, 11, 23.

[9] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier - Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. npj DIGITAL MEDICINE, 3(1), 119.

[10] Quan, Z., Wang, Z., Gan, G., & Valdez, E. A. (2024). Automated machine learning in insurance. arXiv preprint arXiv:2408.14331.