

Improving Business Insurance Loss Models by Leveraging InsurTech Innovation

Zhiyu Quan,¹  Changyue Hu,¹  Panyi Dong,¹  and Emiliano A. Valdez² 

¹Actuarial and Risk Management Sciences, University of Illinois Urbana–Champaign, Urbana, Illinois

²Mathematics, University of Connecticut, Storrs, Connecticut

Recent transformative and disruptive advancements in the insurance industry have embraced various InsurTech innovations. In particular, with the rapid progress in data science and computational capabilities, InsurTech is able to integrate a multitude of emerging data sources, shedding light on opportunities to enhance risk classification and claims management. This article presents a collaborative effort as we combine real-life proprietary insurance claims information together with InsurTech data to enhance the loss model, a fundamental component of insurance companies' risk management. Our study further utilizes a tree-based model and a conventional linear model to quantify the predictive improvement of the InsurTech-enhanced loss model over that of the insurance in-house model. The quantification process provides a deeper understanding of the value of InsurTech innovation and advocates potential risk factors that are unexplored in traditional insurance loss modeling. This study represents a successful undertaking of an academic–industry collaboration, suggesting an inspiring path for future partnerships between industry and academic institutions.

1. INTRODUCTION

InsurTech combines the words “insurance” and “technology” to describe the applications of emerging technology to modernize the entire insurance value chain by improving efficiency, enriching customer service, enhancing underwriting and actuarial processes, and further uncovering new opportunities, to name a few. InsurTech emerged around 2010 and was initially categorized as a subsector of FinTech (Kelley and Wang 2021). InsurTech companies are conventionally funded as start-up companies by venture capitalists; InsurTech reached more than 2000 start-up deals at the end of 2022. The amount of financing has been on the rise globally over the past decade, and in 2021, funds invested totaled \$17.2 billion, breaking records of amounts invested in 2019 and 2020 combined. See Figure 1. The primary investment is directed toward property and casualty (P&C)-focused InsurTech, followed by health, with the lowest allocation to life. Though investment dropped in 2022 due to disruption in supply chains and economic uncertainties affecting inflation and interest rates, InsurTech financing will continue to have a significant presence in the industry for years ahead, with huge potential growth opportunities in the life sector. We continue to see increasing evidence that InsurTech has a significant impact on the local insurance industry in places such as Europe (Ricci and Battaglia 2021), China (Wang 2021), India (Suryavanshi 2022), and Israel (Berman, Shallmo, and Williams 2021).

The progress in InsurTech is accelerated by the rapid increase in the development and adoption of technology, which results in improved connectivity and modern digitization (Fluckiger and Duygun, 2022). Such technological advancements lead to the flow of big data that allows InsurTech innovation to offer user-oriented or data-driven solutions across all major lines of business in the insurance industry, including the automation of business processes, the development of innovative products, and the exploitation of data for underwriting, risk assessment, and claims handling. See VanderLinden et al. (2018) and Xu and Zweifel (2020).

Two of the representative InsurTech applications are telematics in auto insurance and wearable technologies in the life and health sector. For more than a decade, auto insurance companies have used telematics technology to collect information on driving behavior and conditions that would otherwise be unavailable through conventional underwriting (Guillen et al. 2019).

Address correspondence to Zhiyu Quan, Actuarial and Risk Management Sciences, University of Illinois Urbana–Champaign, 270 Computing Applications Building, 1409 W. Green Street, Urbana, Illinois 61801, USA. E-mail: zquan@illinois.edu




FIGURE 1. InsurTech Financing for 2010–2022. *Note:* Data sourced from Financial Technology Partners (2023).

Such information includes, for example, exposure to traffic conditions, driving maneuvering, and distance traveled in real time, which are then used to improve risk assessment, explore the potential impact on pricing, and determine more suitable premiums for usage-based insurance. See Boucher, Côte and Guillen (2017), Weidner, Transchel, and Weidner (2017), Verbelen, Antonio, and Claeskens (2018), Gao, Meng, and Wüthrich (2019), Huang and Meng (2019), and Che, Liebenberg, and Xu (2022). Companies like Discovery Insure, through their Vitality Drive program,¹ exemplify the application of telematics data, rewarding safe driving behavior with lower premiums, fuel cash back, and real-time emergency aid. Root Insurance,² the nation’s first licensed insurance carrier powered entirely by telematics, offers a mobile application to monitor actual driving behavior, allowing for more customized pricing that is suitably affordable for safer drivers but unattractive for reckless drivers. Studies on Root Insurance’s financial reports highlight a noteworthy high combined ratio, prompting further research into the impact of telematics data on insurance claims. Pesantez-Narvaez, Guillen, and Alcañiz (2019) provided extensive numerical analysis on how telematics data can help understand claim occurrences. In the cases where the telematics data cannot be directly associated with insurance claims, Guillen et al. (2020) proposed to introduce near-miss events, which include potentially dangerous behaviors based on certain manually defined criteria, into the telematics data to identify risks. The near-miss telematics suggested by Guillen, Nielsen, and Pérez-Marín (2021) can be utilized as additional ratemaking factors for premium adjustments and incentives to promote safe driving. Furthermore, Pesantez-Narvaez, Guillen, and Alcañiz (2019), Weidner, Transchel, and Weidner (2017), and Gao, Meng, and Wüthrich (2019) emphasized the importance of interpretability of InsurTech variables in the insurance industry and illustrated their models through model coefficients or feature importance. We also witness comparable technological advancements propelled by InsurTech in the health care sector. Wearable technology, or “wearables,” is growing in popularity and is used in the life and health insurance industry to engage customers, collect nontraditional data, and help with underwriting, pricing, loss mitigation, and customer retention (Spender et al. 2019). For example, health insurance companies partner with Vitality Health International,³ empowered by the world’s largest behavioral engagement platform, to track health measurements and reward healthy behavior for customers seeking a healthier lifestyle. McCrea and Farrell (2018) developed a health risk score to investigate how wearable technology could be used to predict overall health and mortality and to improve the pricing of health insurance.

The InsurTech innovations discussed earlier harness emerging technology to provide additional information, empowering insurers to revolutionize conventional insurance models for strategic advantages, but fail to provide empirical evidence on how insurance pricing can be influenced. In this article, we aim to fill this gap by *quantifying* the impact of leveraging InsurTech data to improve loss models, which can then be utilized for ratemaking and risk classification. Consequently, insurers can gain a better understanding of risks within their insurance portfolios, handle claims more efficiently, and further promote more effective marketing. To be more precise, we use data from an InsurTech start-up, Carpe Data,⁴ to examine the efficacy of InsurTech in improving loss models for the entire portfolio of business owner’s policy (BOP) insurance from a mid-sized multistate insurance company. After real-life data analysis and actuarial loss modeling, we are able to demonstrate that

¹<https://www.discovery.co.za/portal/individual/insure-vitality-drive>

²<https://www.joinroot.com>

³<http://www.vitalityhealthinternational.com>

⁴<https://carpe.io>

InsurTech-enhanced models exhibit improved predictability with viable interpretability compared to the existing insurance in-house model. This improvement is attributed to the additional information collected by Carpe Data.

A recent work by Blier-Wong, Lamontagne and Marceau (2024) presented a new approach to home insurance claim prediction by integrating street view imagery using a generalized linear model, employing representation learning to extract advantageous embeddings from the images. However, to the best of our knowledge, there are no published articles in the actuarial literature on InsurTech innovation in business insurance, also known as commercial insurance, for a few reasons. First, Chester et al. (2018) pointed out that the complexity of underwriting and claims, combined with the low volume and the customized nature of transactions, poses obstacles to business insurance that embraces InsurTech. Second, to take advantage of the comprehensive capacity and fully unlock the potential benefits of InsurTech, combining and centralizing data from both insurers and InsurTech is critical, which may pose another challenge for researchers; even in practice, this kind of data sharing is a constraint due to privacy issues. However, we cannot undermine the importance of business insurance. Business insurance premiums are approximately equal to 2.3% of the world's GDP (Organisation for Economic Co-operation and Development, 2023), an indispensable share, which makes the study of business insurance pressing at present.

To address these challenges, we must pinpoint industry partners with extensive experience in business insurance loss models to conduct a comparative analysis. Additionally, collaborating with an InsurTech data vendor that offers solutions specifically designed for business insurance is indispensable. IRisk Lab⁵ at the University of Illinois partnered with Carpe Data and local insurance companies to conduct joint research projects. Industry partners have provided proprietary insurance portfolio information that includes historical loss experience together with insurer in-house rating factors for exposure information, and enriched with InsurTech data. In this article, we present a three-party research collaboration between industry and the university: the InsurTech company, the insurance companies, and the university. This collaboration is unprecedented in actuarial science. We aim to unleash the full potential of InsurTech innovation in improving business insurance loss modeling. Figure 2 demonstrates the flow of information and data in this regard, with IRisk Lab acting as an aggregator to be able to perform data analytics.

The rest of this article is organized as follows. Section 2 provides a brief description of business insurance, exploratory data analysis, and a mathematical justification for the advantages of bringing InsurTech innovation to the loss model. Section 3 describes the calibration of the model and the results. Section 4 provides model interpretation and suggestions for potential rating factors. We conclude in Section 5.

2. OPPORTUNITIES FOR INSURTECH IN BUSINESS INSURANCE

Business insurance offers coverage designed to protect businesses and organizations of various sizes and industries from diverse risks and financial losses. It encompasses property insurance, liability insurance, business interruption insurance, commercial auto insurance, workers' compensation, and more. Additionally, bundled insurance options such as BOP and commercial multiperil (CMP) are available within the spectrum of business insurance.

BOP aims to protect small and midsize business owners from potential business risks through a bundle of insurance policies with various coverages. See Naylor (2017). Most BOP bundles contain essential coverage of property insurance and liability insurance, and additional coverage, such as commercial auto insurance, commercial crime insurance, business interruption insurance, and other business insurance policies, can be included through customization or negotiation. CMP is often more suitable for larger businesses with complex insurance requirements, given its flexibility and ability to address a wide range of risks. The underwriting procedure for business insurance necessitates a more thorough analysis and assessment, leading to increased underwriting expenses. Delivering a heightened level of customization and personalized service could incur higher administrative costs. Additionally, the distribution channels for business insurance often include more intermediaries, such as agents and brokers, which affects the expense ratio. We select three representative insurance companies, namely, Travelers for large size, Hanover for midsize, and Donegal for small size, offering a diverse view of insurers based on their scale. Their respective market shares and detailed loss summaries, extracted from their annual reports, are presented in Table 1. Here, Δ represents the difference between the operational ratios (loss ratio, expense ratio, combined ratio) of the selected companies and the average ratios in the P&C sector. It is evident from the table that business insurance exhibits a higher expense ratio, approximately 4% to 5%, in contrast to personal insurance. The business insurance expense ratio is around 30%, almost half the corresponding loss ratio. Furthermore, smaller insurers display a higher expense ratio, especially in the context of business insurance, where the difference is more pronounced compared to larger insurers. Despite human intervention in the

⁵See IRisk Lab currently serves as an academic-industry collaboration hub, facilitates the integration of discovery-based learning experiences for students, and showcases state-of-the-art research in all areas of risk analysis and advanced analytics.

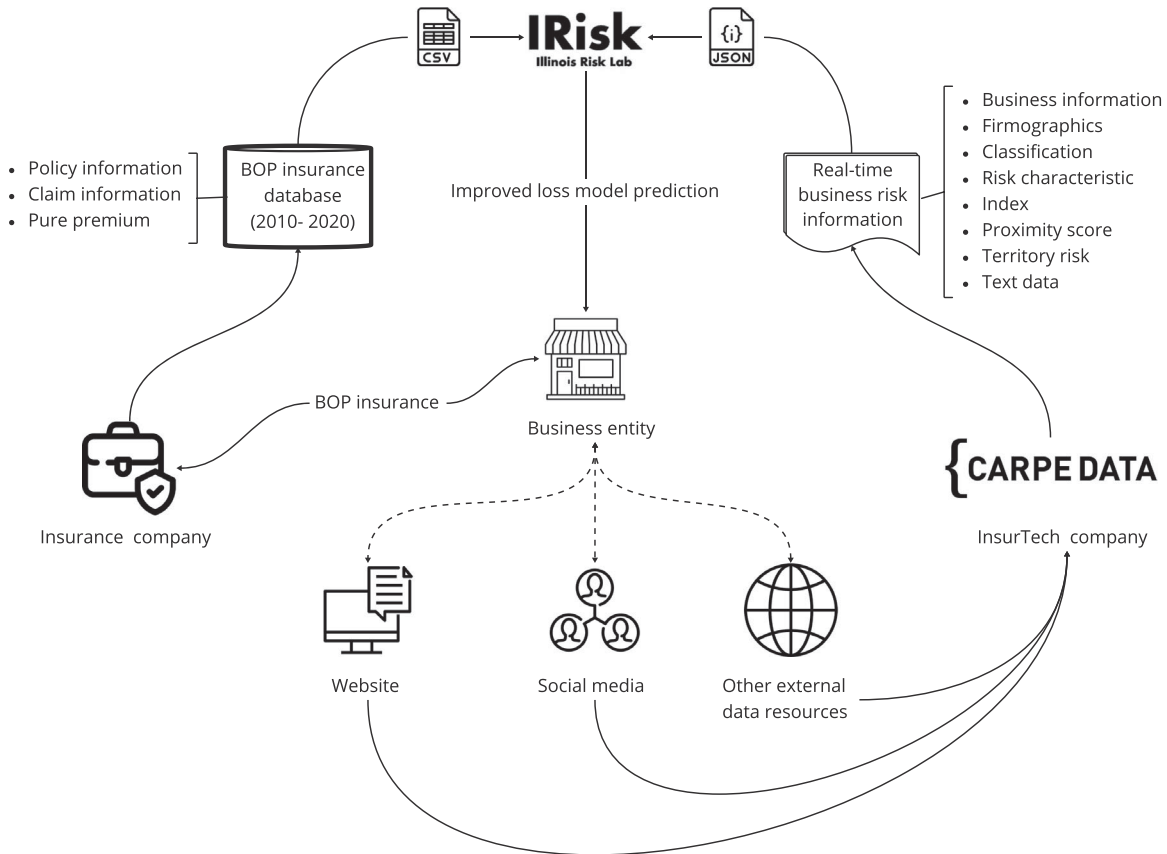


FIGURE 2. The Flow of Information from the Academic–Industry Collaboration.

underwriting process, insurers fall short of gathering sufficient risk factors essential for accurate pricing, resulting in elevated loss ratios. This trend is commonly observed in bundled business insurance across the industry. Regardless of the size of the insurers and their superior performance in general business insurance compared to the baseline, all selected companies show that for bundled business insurance such as CMP, the reserve for claims and loss adjustment expense (LAE) is higher than the written premium. In the case of our insurer partner, we observe the actual/expected ratio of 107.52%, representing the ratio of actual losses to those predicted by the insurer’s in-house loss modeling, for the entire BOP bundle.

The current circumstances clearly require improvements, and there is a demand for effective solutions. Fortunately, InsurTech can provide a customized underwriting system with dynamic programmable rules. The InsurTech-enhanced features, as another typical InsurTech application, empower underwriters and claims adjusters with access to enhanced information about policyholders and claims, facilitating more informed decision making. This is one of the reasons insurance companies explore alternative solutions from InsurTech, and we observe active collaborations between insurance companies and InsurTech companies conducting pilot studies to assess cost–benefit in the present day. For instance, a customized underwriting system can automatically prefill information for the standard underwriting process, saving time for underwriters that would otherwise be spent on communicating with policyholders or manually searching for information from various sources. This efficiency has the potential to drive costs down and even enable process automation for smaller policies such as BOP without the need for direct intervention from underwriters. By harnessing new and multiple external data sources, including social media data and online content, analytics are empowered and accurate information is gained on risks and policyholder behavior. Furthermore, data enrichment and improved internal data management contribute to more precise pricing and the formulation of data-driven underwriting strategies for traditional insurance companies.

In the subsequent section, we illustrate our datasets gathered from industry partners; an anonymous insurance company shares its proprietary loss experience over the past decade, which is an entire database created to make current rate plans for BOP, along with partial rating factors for in-house loss modeling. The InsurTech company, Carpe Data, provides dynamic and real-time information from emerging public data sources that illuminate every aspect of a business, including operations, products, services, and reviews. It offers potential risk factors by connecting the different risk profiles. In summary, the response

TABLE 1
Earned Premium and Claim Payout Information by Insurers

Company	Business line	Information ^a	Year		
			2022	2021	2020
Travelers ^a	CMP	Written premium (in millions)	4,304	3,768	3,608
		Reserve for claims and LAE (in millions)	5,383	4,977	4,665
	Business	Earned premium (in millions)	17,095	15,734	15,294
		Δ Loss ratio (%)	-13.6	-7.5	-0.7
		Δ Expense ratio (%)	3.8	4.2	3.4
	Personal	Δ Combined ratio (%)	-10.2	-4.0	1.5
		Earned premium (in millions)	13,250	11,983	10,927
		Δ Loss ratio (%)	3.4	-2.2	-7.3
	Hanover ^b	CMP	Δ Expense ratio (%)	-0.8	-0.3
Δ Combined ratio (%)			2.2	-3.2	-9.1
Business		Written premium (in millions)	1,011	979	922
		Reserve for claims and LAE (in millions)	1,557	1,338	1,184
Personal		Earned premium (in millions)	1,951	1,811	1,704
		Δ Loss ratio (%)	-7.9	-5.1	-6.9
		Δ Expense ratio (%)	6.8	6.4	5.7
Personal		Δ Combined ratio (%)	-1.5	0.6	-2.4
		Earned premium (in millions)	2,113	1,929	1,844
	Δ Loss ratio (%)	1.4	-4.0	-5.4	
Donegal ^c	CMP	Δ Expense ratio (%)	0.6	1.2	0.2
		Δ Combined ratio (%)	1.6	-3.5	-6.4
	Business	Written premium (in millions)	200	188	148
		Combined ratio (%)	116.9	114.1	98.4
	Personal	Earned premium (in millions)	510	468	413
		Δ Loss ratio (%)	-9.3	-3.9	-6.2
		Δ Expense ratio (%)	9.6	9.8	6.0
	Personal	Δ Combined ratio (%)	1.0	5.2	-1.0
		Earned premium (in millions)	312	307	329
Δ Loss ratio (%)		-5.4	-7.7	-10.6	
Baseline ^d	P&C	Δ Expense ratio (%)	5.9	2.1	4.8
		Δ Combined ratio (%)	0.1	-5.3	-3.4
		Loss ratio (%)	76.4	72.5	70.1
		Expense ratio (%)	25.9	26.5	27.5
		Combined ratio (%)	102.7	99.7	98.8

Notes: Δ Ratio = Company ratio – Baseline ratio. ^aData from The Travelers Companies (2022). ^bData from The Hanover Insurance Group (2022). ^cData from The Donegal Insurance Group (2022). ^dData from Federal Insurance Office, U.S. Department of the Treasury (2023).

variable is the BOP loss cost during the observation period. For predictive features, we have two sources: the insurance company provides policy information used as rating factors and Carpe Data offers policyholders' risk factors.

2.1. Insurer's BOP Policy Information

In our BOP dataset, there are three types of coverage: business building (BG), business personal property (BP), and liability (LIAB). BG explicitly covers the loss related to the buildings in which the business operates, owned or rented by the business owners. BP covers the risks of potential loss, damage, and liability issues for the property of policyholders that has a business use. The losses covered by BG and BP are discovered and covered quickly (short-tail) with relatively minor claims, which are often capped by the value of the properties. LIAB covers the risks of potential legal liability for property damage, bodily

TABLE 2
Policy Information and Loss Experience Provided by the Insurance Company

Category	Features	Details
Policy information	Policy year	From 2010 to 2020
	Earned exposure	Unit of exposure in years
	Coverage limit	The maximum amount that the insurance company pays for a covered loss, in dollars
	Coverage type	BG, BP, LIAB
	Exposure base	Limit of Insurance (LOI), annual gross sales, annual payroll
	Risk type	Apartment, condo/office, contractors, convenience, distributor, fast food, motel, office, other, restaurant, retail, self-service
Loss experience	Observed loss cost	Actual loss cost observed, in dollars
	Insurance company's in-house model loss cost	Predicted loss cost by the insurance company's in-house model, in dollars

injury, and other losses caused by policyholders to a third party; for example, medical payments to customers who are injured while on business property, defense costs, damages for failure to or improperly rendering professional service, and accidental pollution. Compared to BG and BP, the losses covered by LIAB are usually long-tail risks, where the report and settlement can take a long time, and the frequency and severity of claims vary depending on industry and type of incident.

As mentioned earlier, the response variable is the BOP loss cost observed from 2010 to 2020. Table 2 outlines the response variable together with the features provided by the insurance company. It is worth mentioning that the insurance company also attaches a pure premium for each policyholder, which is the insurance company's in-house model loss cost. Due to privacy concerns, we do not have all the rating factors that are used in the in-house model. However, this does not preclude us from directly comparing our results with the in-house model results, which will be discussed in Section 4.

Figure 3 shows the distribution of our response variable or the observed loss cost by the three coverage groups. The violin plots show the kernel density for the logarithm of the observed loss cost plus one. We can observe that the majority of policies have zero loss costs, meaning that most policyholders do not have claims submitted. Based on the shapes of the distributions, we see that the loss cost of the LIAB group is centered around zero with the most significant proportion. Specifically, 96.50% of the policies of the BG group, 98.99% of the policies of the BP group, and 99.20% of the policies of the LIAB group have zero claims. In addition, we present histograms showing how the logarithm of the positive observed loss cost plus one is distributed. We can deduce that the median value of the observed loss from the LIAB group (around e^{12}) is higher than that of the other two groups. The other two groups, BG and BP, are around e^9 and e^{10} , respectively. Finally, we observe that some policyholders tend to report claims more often or to report more expensive claims than others. In practice, most insurance portfolios are heterogeneous because they are usually a mix of policyholders with different risk levels. This is especially true for BOP with three different coverages. After basic exploratory data analysis and comprehensive model experiments, we decided to model these three coverage groups separately.

Next, we focus on exploratory data analysis and get a first impression of how features influence the response variable. Clearly, due to space constraints, we can show only some of the representative features. However, we focus on some of the features that appear to be worth discussing after model fitting.

Let us pick one rating factor, the risk type, from the policy information provided by the insurance company to show one-way analyses. We note that in this and subsequent preliminary analysis, the observed loss cost and features analyzed are for BG coverage only, due to the limitation of article length. Figure 4 demonstrates eleven risk types in total, ranging from *Apartment* to *Self-Service*. Figure 4(a) shows the number of policies based on the different types of risk. We can see that the *apartment* risk type has the most significant amount, followed by the *office* risk type and the *retail* risk type, whereas the *convenience* risk type has the least amount. Figure 4(b) depicts the average value of the observed loss cost for each different type of risk. The black line shows the standard deviation of the observed loss cost. On average, the *condo/office* risk type BOP policies have the highest observed loss cost, whereas those of the *self-service* risk type have the lowest. Figure 4(c) presents the violin plots of the logarithm of the observed loss cost plus one by the different types of risk. The shapes of the distributions reveal that the loss cost from the *convenience* risk type is the most dispersed, whereas the *self-service* risk type is the most concentrated. To highlight the differences in the observed loss cost by different types of risk, we also show box plots of the

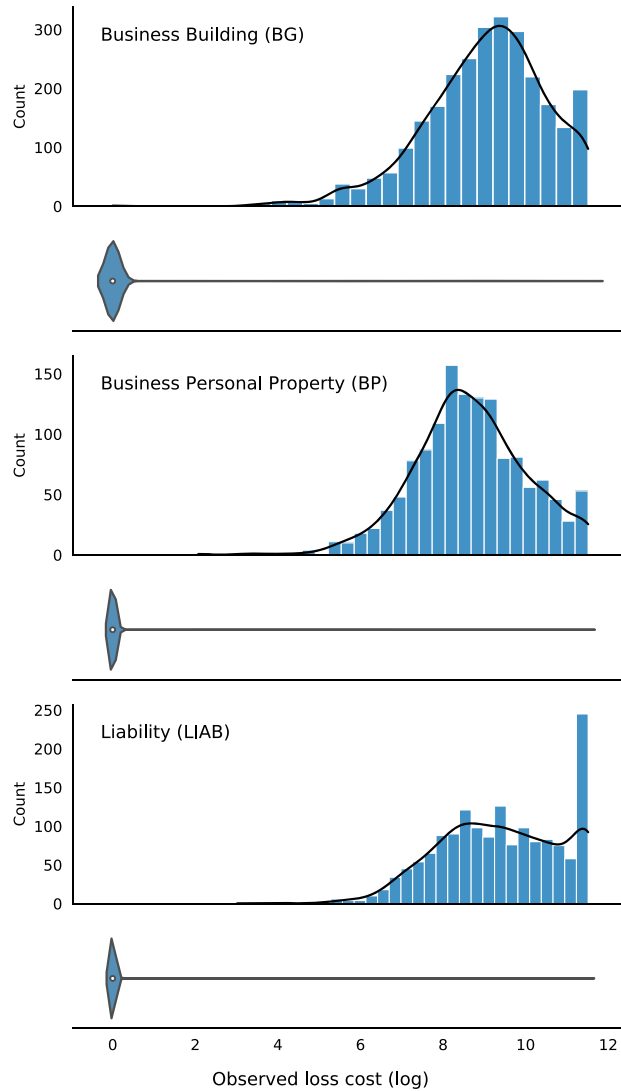


FIGURE 3. Distribution of the Observed Loss Cost by Coverage Type.

logarithm of the positive observed loss cost in Figure 4(d). It shows that the median values across all types of risk are relatively similar, floating around 8000 (e^9); however, the dispersion of the different types of risk varies significantly. This is yet additional evidence of the complexity of BOP across multiple industries and under extensive risk exposure.

2.2. Data from InsurTech

The potential of InsurTech innovation to help the insurance company model these extensive risks has never been greater. However, unlike most academic toy datasets, we cannot safely assume that our real-life dataset is reliable and accurate. First, we need to carefully check the entity resolution, which evaluates whether the InsurTech company is able to match the correct policyholder based on the limited information, including address, business name, and others. Second, we need to perform data engineering work, including extracting structured datasets from the raw data files provided in nested JSON (JavaScript Object Notation) format with hierarchical structures; cleaning dirty data, which includes missing data, inconsistent format, and human input errors; and creating useful features from various data types, including text data. Third, we need to set up our relational database that connects the InsurTech features derived from various data sources to each policyholder of the insurance

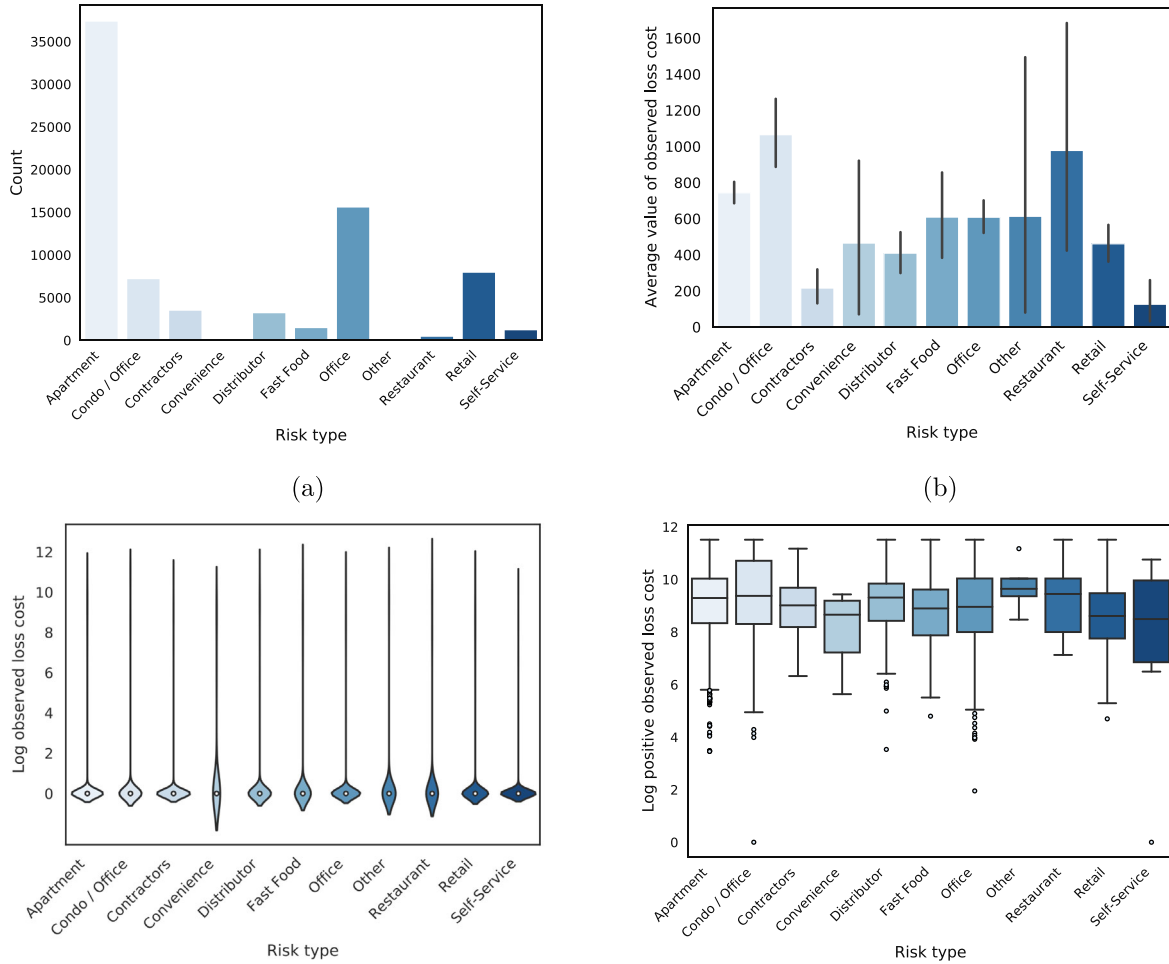


FIGURE 4. Risk Type: (a) Bar Plot of Risk Type, (b) Bar Plot of Average Observed Loss Cost by Risk Type, (c) Violin Plot of Log Value of Observed Loss Cost by Risk Type, and (d) Box Plot of Log Value of Positive Observed Loss Cost by Risk Type.

company. The final version of the combined dataset has 825,622 observations and 596 features. Table 3 broadly summarizes various categories of InsurTech features with explanations and examples.

The InsurTech features of the classification category have information similar to the risk types mentioned in Section 2.1 but provide a more granular categorization of a business. There are 24 segments that provide categorization of a business, varying from *accountants & financial services* to *wholesale & distributors*. Additionally, companies can operate in multiple lines of business, thus falling into multiple segments. Carpe Data utilizes various classifiers to assign a segment to a business, and segment details provide the proportion of total votes of a certain segment for a business. For example, one of the policyholders is a grocery store run by a family farm that sells and delivers meat to the surrounding community. It belongs to the *retail* risk type based on insurance policy information and is also segmented as *retail* (with a majority vote of 50%) from the InsurTech aspect. Furthermore, based on segment details information, it is also categorized as *wholesale & distributors* (20%, sells and delivers meat to the surrounding company), *agriculture, forestry, fishing & hunting* (10%, farm), *food services & drinking places* (10%, grocery has food service), and *manufacturing* (10%, produce dairy products). Despite the overlap between risk types from in-house rating factors and segment information from InsurTech innovations, InsurTech features can cross-validate and enhance internal rating factors.

We now focus on exploratory data analysis for some of the representative InsurTech features mentioned in Table 3. Indexes are a suite of numerical features that measure the risks faced by a business. Among them, the visibility index measures a business entity's visits and online presence based on social media, visitor reviews, and marketing information. Figure 5 illustrates the visibility index and its univariate relationship with the observed loss cost. Figure 5(a) shows the distribution of the visibility index, which ranges from 1.0 to 5.0 and peaks at about 3.2. This score helps measure the absolute risk of a business, the

TABLE 3
Features from InsurTech Company

Category	Explanation	Examples
Business information	Basic information concerning the operation of businesses including address, year founded, operating hours, etc.	<i>coordinates, state, street_type, type_of_address, is_franchise, is_home_business, operating_year, language_spoken, opening hours</i>
Firmographics	Characteristics to segment prospect business including the type of business, the number of employees, the sales ranges, etc.	<i>business_size, credit_score, company_type, sales_range, revenue_range</i>
Classification	Categorization of a business	<i>category, segment, segment_details, NAICS code</i>
Risk characteristics	Indicator features that identify current attributes relating to the potential risks of a business	<i>Alcohol, chemical application, fireplace, gasoline, heavy construction, outdoor heaters, raw seafood served, tobacco, work at heights</i>
Index	A suite of numerical indexes on a 1 to 5 scale targeting dimensions of risk that can be tuned by segment and location	<i>Customer rating, visibility, reputation, health & sanitation, maintenance & condition</i>
Proximity score	Proximity scores identify risks associated with the surrounding businesses that may impose on policyholders in the immediate vicinity of that risk	<i>Negative keywords, combustibles, entertainment, traffic</i>
Territory risk	Density scores of risks within a ZIP code area, engineered from collected information based on the location of a business	<i>TERRITORY.a1 to TERRITORY.m1 TERRITORY.a2 to TERRITORY.m2</i>
Text data	(1) Web page content crawled from a business's website and (2) customer reviews of a business from multiple data sources, including but not limited to the content of the review, star rating, number of likes, the date of posting, and the source of the reviews.	<i>Web page content: content, title, url Customer reviews: content, likes, dislikes, stars, saves, language</i>

risk relative to other businesses, and the change in risk over time. [Figure 5\(b\)](#) presents a scatterplot of the visibility index and the positive observed loss cost with a nonlinear regression line. The trend of the regression line shows that the loss costs are higher when the visibility index is lower and vice versa. Visibility is an important tool for identifying potential customers to improve business performance.

Proximity scores identify risks associated with surrounding businesses that can expose policyholders to additional risks in their immediate vicinity. Three groups of proximity scores are considered in our datasets: combustibles, entertainment, and traffic. Each group has two scores: the density score and the distance score. The density score is calculated over a certain range by counting the total number of risks and normalized to a value that is an integer from 1 to 3, where 1 indicates high risk and 3 indicates low risk; the distance score measures the nearby risks as a function of distance. It is a float number between 1.0 and 5.0, with a higher number indicating a lower degree of risk. [Figure 6](#) illustrates one group of proximity scores, which is proximity traffic, and shows how it relates to the observed loss cost. As we mentioned, there are two types of scores for proximity traffic, the density score and the distance score. [Figure 6\(a\)](#) compares the number of policies with different density scores, showing that the number of policies increases with the density score, with a density score of 3 having the highest number of policies. Similarly, [Figure 6\(c\)](#) illustrates a more granular view of how distance scores are distributed, and we observe a consistent trend that there are more observations with a distance score close to 5. [Figure 6\(b\)](#) shows that there has been a steady decline in the average value of the observed loss cost as the density score of proximity traffic increases, which matches the intuition that a lower surrounding risk implies a lower claim loss. Additionally, businesses with a density score of 1 have the highest variability in the observed loss cost. [Figure 6\(d\)](#) presents a scatterplot of the distance score and the observed loss cost

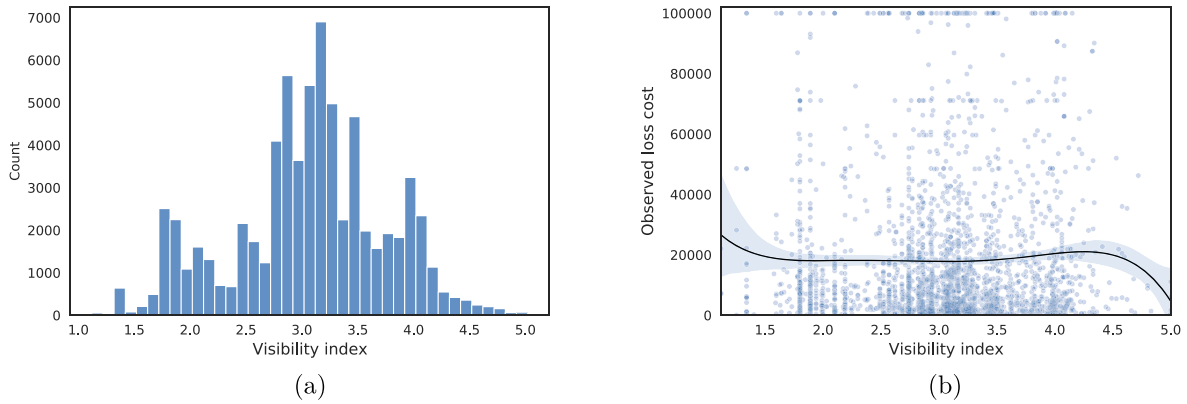


FIGURE 5. Visibility Index: (a) Histogram of Visibility Index and (b) Scatterplot of Visibility Index and Observed Loss Cost with a Nonlinear Regression Line.




FIGURE 6. Proximity Traffic Scores: (a) Bar Plot of Traffic Density Score, (b) Bar Plot of Average Value of Observed Loss Cost by Traffic Density Score, (c) Histogram of Traffic Distance Score, and (d) Scatterplot of Traffic Distance Score and Observed Loss Cost with a Nonlinear Regression Line.

with a nonlinear regression line to get a general sense of the relationship. We see that this nonlinear regression line has a nearly constant trend, but the confidence interval becomes narrower with more observations, especially as the proximity traffic distance score reaches 5. Usually considered as a feature that may indicate accessibility, proximity to traffic can be a complex variable with an unclear impact on business performance, because its effect can vary widely depending on the type of business.

Territory risk measures the cluster of risks within the ZIP code level area, and it is engineered from the business risk characteristics defined in Table 3. Because it is determined based on the ZIP code level, it can be readily incorporated with existing location-based rating factors for insurance company use. There are 13 clusters of risks, named risk A to risk M, that focus on

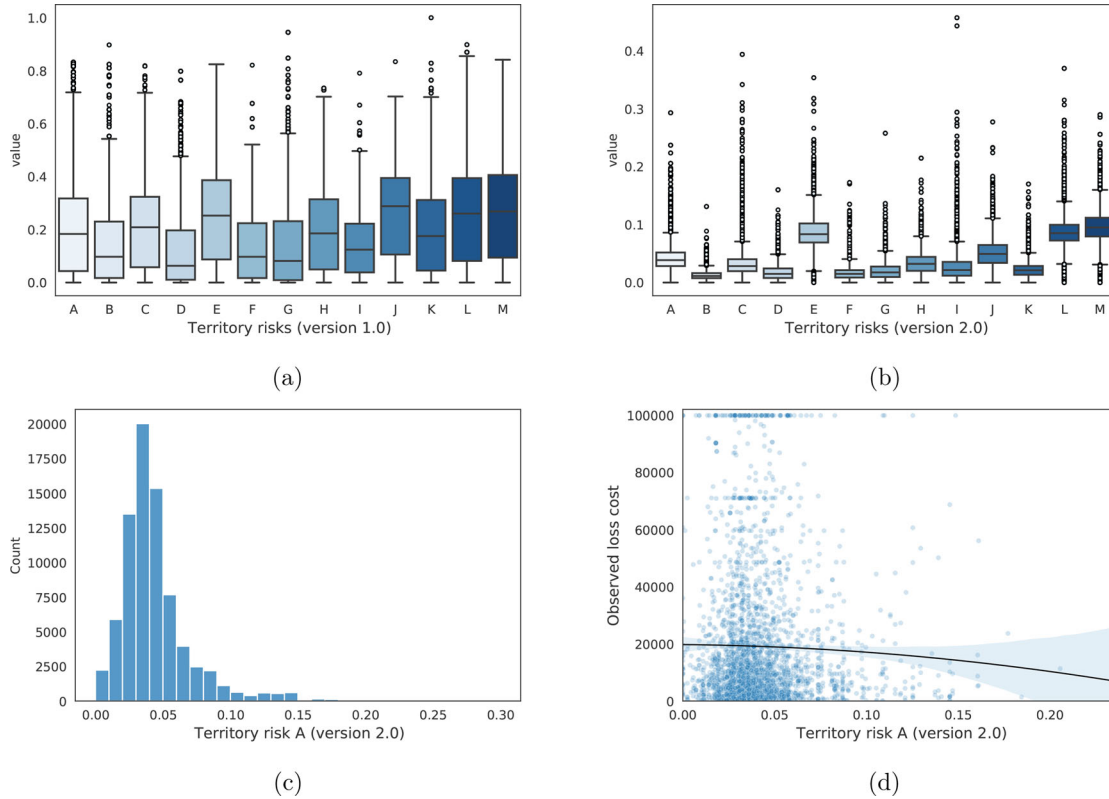


FIGURE 7. Territory Risk: (a) Box Plot of Territory Risk (Version 1.0), (b) Box Plot of Territory Risk (Version 2.0), (c) Histogram of Territory Risk A (Version 2.0), and (d) Scatterplot of Territory Risk A (Version 2.0) and Observed Loss Cost with a Nonlinear Regression Line.

different aspects of the business risk characteristics. For example, one of the clusters may focus on fire-related risk characteristics concerning potential fire hazards in this territory. Based on different feature engineering techniques, two versions of territorial risk, namely, version 1.0 and version 2.0, are used in our modeling. Both versions are continuous and range between 0 and 1 for version 1.0 and between 0 and 0.5 for version 2.0. For each version, a higher value indicates a lower risk. Figure 7(a) and Figure 7(b) show the box plots of territorial risk with two different versions. We can see that the features of version 1.0 are more dispersed than those of version 2.0. Figure 7(c) shows the right-skewed distribution of territorial risk feature A, ranging from 0 to 3.0, with the most common value of 0.04. Figure 7(d) presents a scatterplot of the territorial risk feature A (version 2.0) and the observed loss cost.

2.3. Information Gained from InsurTech

Based on datasets obtained from our industry partners, we focus on improving BOP loss models by leveraging InsurTech innovations. P&C insurance carriers started to apply modern predictive modeling in the early 2000s. Meanwhile, insurance companies began to adopt analytics with large insurance datasets for various applications. For example, actuarial ratemaking, or actuarial pricing, described in Denuit et al. (2007), refers to the process during which actuaries quantify the insurance products for the underlying risks covered by insurance policies. The process of ratemaking targets to accurately predict future claim losses by policyholders based on historical loss data. One of the most significant underpinnings of the ratemaking process is data. The accuracy and fairness of the premium rates that are ultimately determined depend largely on the quality and quantity of available data. Therefore, it is natural to look for good-quality external data to enhance the underlying predictive models. In addition, it is worth mentioning that insurance is highly regulated, and business insurance premiums must comply with the laws set by local regulators. Such regulation usually serves two purposes: first, to ensure the insurer's solvency and, second, to protect customers by restricting unfair rating factors. Hence, in this article, our goal is twofold: mining predictive risk characteristics from InsurTech data that improve insurance in-house loss model and explaining these risk characteristics using interpretable machine learning techniques. We also try to further propose potential rating factors for business insurance.

Let us define \mathbf{X}^{IH} as in-house rating factors used by the insurance company, such as policy information, basic policyholder information, and conventional risk characteristics. \mathbf{X}^{IT} denotes risk factors generated by InsurTech innovations, such as enhanced policyholder information, firmographics and territory information, and related social media information. We assume that \mathbf{X}^{IH} and \mathbf{X}^{IT} are subsets of \mathbf{X}^W , which is a set of ground-truth risk factors that perfectly define the underlying risk. Obviously, \mathbf{X}^W is unobtainable and unrealistic in a real-life setting. We further assume that there is no overlap between \mathbf{X}^{IH} and \mathbf{X}^{IT} without loss of generality. Indeed, there may be some overlap in the information between insurance in-house rating factors and InsurTech risk factors. However, the choice to utilize overlapping data usually entails a careful consideration of the trade-offs between costs and benefits. Such overlapping information can serve as a cross-reference to validate insurance in-house rating factors, potentially paving the way for further improvements.

Let Y be the claim amount, which is our response variable, and the variance of Y can be decomposed according to the law of total variance as

$$\text{Var}(Y) = \text{E}[\text{Var}(Y|\mathbf{X}^{IH})] + \text{Var}(\text{E}[Y|\mathbf{X}^{IH}]), \quad (1)$$

where $\text{E}[Y|\mathbf{X}^{IH}]$ is the current pure premium based on the insurance in-house loss model. The insurer believes that given current rating factors, $\text{Var}(Y|\mathbf{X}^{IH})$ is considered random as white noise, because the variability in claim amounts is solely due to chance. However, this belief can be arguable, especially when there is a large information discrepancy between \mathbf{X}^{IH} and \mathbf{X}^W . The insurer must carefully monitor this variance in the presence of systematic drift. The second term can be deciphered from the portfolio perspective; that is, based on the current pure premium, $\text{E}[Y|\mathbf{X}^{IH}]$, it is the variance of the current pure premium that depends on the portfolio composition. If the current pure premium is not precise enough, then adverse selection could occur, and the change in portfolio composition could generate systematic losses for the insurer.

The first term can be further decomposed, using Bowsher and Swain (2012), by bringing InsurTech risk factors, \mathbf{X}^{IT} , as follows:

$$\text{Var}(Y) = \text{E}[\text{Var}(Y|\mathbf{X}^{IH}, \mathbf{X}^{IT})] + \text{E}[\text{Var}(\text{E}[Y|\mathbf{X}^{IH}, \mathbf{X}^{IT}]|\mathbf{X}^{IH})] + \text{Var}(\text{E}[Y|\mathbf{X}^{IH}]), \quad (2)$$

where $\text{E}[Y|\mathbf{X}^{IH}, \mathbf{X}^{IT}]$ is the updated pure premium by leveraging InsurTech innovation. Evidently, $\text{E}[\text{Var}(Y|\mathbf{X}^{IH}, \mathbf{X}^{IT})]$ is smaller than $\text{E}[\text{Var}(Y|\mathbf{X}^{IH})]$ and could be closer to random noise. The second term explains the variation derived from the imperfect risk classification. More specifically, this improved pure premium, or the refined risk classification based on InsurTech innovation, $\text{E}[Y|\mathbf{X}^{IH}, \mathbf{X}^{IT}]$, cannot be fully explained by the current rating factors \mathbf{X}^{IH} alone. As a result, incorporating additional InsurTech risk factors extends several benefits in insurance ratemaking. Expanded risk factors can help develop a better risk classification and attract more homogeneous policyholders. Hidden risk factors may be uncovered, thereby avoiding premium leakage that can lead to systematic losses. A more comprehensive understanding of risk factors can help the insurer gain relative advantage over its competitors and further control the variance of the loss.

In addition, we can adopt the credibility theory to further justify the advantages of bringing in external information in the future. Frees and Shi (2017) addressed how individual insurers blend external information with their own rating variables to improve the ratemaking process by developing credibility predictions within a generalized linear model framework, with a focus on the Tweedie family, using Bayesian methods. Diao and Weng (2019) proposed a regression tree credibility model, incorporating Bühlmann-Straub credibility loss functions into regression tree models, leading to a higher prediction accuracy of credibility premiums.

3. MODEL CALIBRATION

To date, a wide range of actuarial pricing theories, strategies, and properties have been proposed, ranging from classical statistical methods to modern machine learning techniques such as deep learning approaches (e.g., Denuit, Charpentier, and Trufin 2021; Chen et al. 2023; and Zhu 2024). We investigated several algorithms to examine the effectiveness of InsurTech innovation through data enhancement, many of which were unsuitable for our purpose. We meticulously chose to present three representative methodologies related to loss modeling: Tweedie generalized linear model (GLM) in Appendix A.1, elastic net regularization in Appendix A.2, and light gradient-boosting machine (LightGBM). These three methodologies have been used to find a function f that maps $(\mathbf{X}^{IH}, \mathbf{X}^{IT}) \xrightarrow{f} \text{E}[Y|\mathbf{X}^{IH}, \mathbf{X}^{IT}]$. In other words, supervised learning models predict a response variable, pure premium, from a function, f , of risk factors and parameters.

3.1. Light Gradient Boosting Machine

Actuaries nowadays resort to advanced machine learning methods to accurately assess policyholders' risk profiles. Machine learning methods provide state-of-the-art performance and high efficiency. Wuthrich and Buser (2021) summarized some of these methods for actuarial ratemaking and compared them with traditional actuarial models. We also observe that advanced machine learning methods, including those we considered outside the reach of this article, can significantly outperform the classical Tweedie GLM using our datasets.

Though machine learning methods are acknowledged for their advanced predictive capabilities, concerns have been mounting regarding their interpretability, especially in light of regulations such as the General Data Protection Regulation in Europe. These regulations emphasize individuals' right to explanation, a crucial consideration in industries like insurance where transparent pricing methods are not only advisable but also mandated by regulators. GBMs, beyond offering competitive predictive performance in our dataset, provide interpretability, offering insights into the model's decision-making process. Though each individual tree within a GBM is inherently explainable, the collective result of numerous trees can be intricate. However, as highlighted in Delgado-Panadero et al. (2022), techniques have been developed to extract feature contributions from GBMs, thereby shedding light on prediction formulations. Such techniques enhance the interpretability of GBMs and align well with regulatory mandates like General Data Protection Regulation, effectively balancing predictive accuracy with model explainability. With these merits of tree-based models in mind, we now delve into their applications in the actuarial science literature.

Guelman (2012) applied GBMs to the prediction of auto at-fault accident losses and demonstrated that it produces a superior predictive accuracy than the traditional GLM approach. Quan and Valdez (2018) modeled multi-line of business insurance claims data with correlated responses using multivariate tree-based models and showed that multivariate tree-based models outperform univariate ones in terms of predictive accuracy and are able to capture the inherent associations among response variables. Lee and Lin (2018) introduced delta boosting and demonstrated that it is the most optimal boosting method for various loss functions and tested the proposed delta boosting machine with collision insurance claim data. Yang, Qian, and Zou (2018) proposed TDBOOST, a gradient tree-boosted Tweedie compound Poisson model, and applied the proposed method to the prediction of auto insurance claims. As an extension of TDBOOST, Zhou, Qian, and Yang (2022) proposed EMTBOOST, a gradient tree-boosted zero-inflated Tweedie model, for extremely imbalanced data with a large number of zeros. Henckaerts et al. (2021) implemented tree-based models with tailored loss functions for the development of full tariff plans based on both the frequency and severity of claims and concluded that the boosted trees outperform the classical GLM approach. Hu, Quan, and Chong (2022) modified the traditional splitting function of decision trees to handle the imbalance problem induced by a large proportion of zero responses in insurance claim datasets and demonstrated the modified tree structure leads to improved prediction accuracy. Qiao, Wang, and Zhu (2024) proposed an ensemble modeling framework that leverages various tree boosting models for long-term mortality forecasting. Their empirical analysis indicates that this method surpasses both the classic stochastic mortality models and deep learning-based benchmarks in terms of predictive accuracy.

More specifically, LightGBM, which is satisfactory for our purpose of seeking an efficient implementation of gradient boosting decision tree (GBDT), is an ensemble of decision trees (DTs) in Appendix A.3, where each proceeding tree is grown sequentially based on the negative gradient or residual of previously learned decision trees.

To extend a single DT to GBDT, a series of trees are grown in a sequential fashion, where every tree is targeted to optimize the negative gradient of previous trees. As indicated in Friedman (2001), the true function f can be approximated by

$$f^B(\mathbf{X}) = \sum_{b=0}^B f_b(\mathbf{X}), \quad (3)$$

where B denotes the number of functions or iterations, f_0 is the initialization function, and $\{f_b\}_{b=1}^B$ are the incremental functions. In the case of GBDT, each f_b is a DT and f_b is grown at iteration b . In the circumstance of steepest descent with classification and regression tree, if we define the sum of trees up until the current step B as $f^B(\mathbf{X}) = \sum_{b=0}^B f_b(\mathbf{X})$, the tree update procedure can be formulated as

$$f^B(\mathbf{X}) = f^{B-1}(\mathbf{X}) + \rho_B \sum_{m=1}^{M(B)} \hat{c}_{mB} \mathbb{1}_{R_{mB}}(\mathbf{X}), \quad (4)$$

where ρ_B is the predefined learning rate, $M(B)$ refers to the number of leaf nodes at iteration B , R_{mB} is the leaf node m at iteration B , and \hat{c}_{mB} is the prediction constant of node R_{mB} , which is estimated by averaging the node, as stated above.

Furthermore, as stated in Ke et al. (2017), to reduce the computational complexity induced by the recursive binary splitting process, where each internal node needs to scan every feature and every observation, LightGBM innovatively utilizes gradient-based one-side sampling (GOSS) and exclusive feature bundling (EFB) to address the computation issues. GOSS is a sampling mechanism that assigns higher weights to observations with larger gradients, which leads to more contributions to the incremental functions. It also employs a down-sampling technique to retain those observations with a smaller gradient, which leads to randomly dropping less influential observations before the tree construction procedure. This method is typically better than random sampling, which assigns the same sampling rate. Furthermore, to reduce the dimension of the feature space, especially when we have a sparse dataset, EFB proposes to combine the exclusive features as bundles whose size is significantly reduced compared to the original feature space. Through GOSS and EFB, Ke et al. (2017) claimed to substantially speed up GBDT while maintaining comparable accuracy, which is one of the reasons we selected this algorithm given that we have a large dataset with almost 1 million observations and more than 500 features.

All machine learning methods need proper model calibration, which involves a train–test split, cross-validation, and hyperparameter tuning. It requires massive computation and various experiments. We are able to leverage the HAL system⁶ to perform cloud computing in a secured environment. See Kindratenko et al. (2020). In addition, we use the Optuna framework to perform hyperparameter optimization, specifically using Bayesian optimization. See Akiba et al. (2019) and Shahriari et al. (2015). The Optuna framework is able to fully utilize distributed cloud computing by allowing users to submit multiple batch jobs on the same training process, which narrows down the optimal hyperparameter space with the Bayesian optimization mechanism.

3.2. Comparison of Model Performance

This section compares the performance of these two models: insurance in-house models versus InsurTech-enhanced models. As we mentioned in Subsection 2.1, we model each coverage group separately using LightGBM and Tweedie GLM after elastic net feature selection, and thus we have six InsurTech-enhanced models in total. For LightGBM model calibration, we use Bayesian optimization with the Optuna framework, as discussed in Section 3.1, and the optimization objective loss function is the mean absolute error (MAE). For the Tweedie GLM after elastic net feature selection model calibration, we use grid search to find the optimal hyperparameters mentioned in Appendix D. We apply a 10-fold cross-validation on the training set to find the best models for both methods. For the hyperparameter values of the final models after tuning, see Appendix B.

We use double lift charts to visually compare the predictive performance of InsurTech-enhanced models and insurance in-house models. Model lift refers to the ability to differentiate between low- and high-risk policyholders and can be used to measure a model’s business value, which helps c-suite leaders make a business decision. Double lift charts are commonly used in insurance companies to measure model lift and compare predictive power between two distinct models. In our experiments, double lift charts are created according to the following procedures: First, we sort the data by a ratio of new model prediction (InsurTech-enhanced model prediction) to the current pure premium (insurance in-house model prediction). Second, the sorted data are subdivided into quantiles with equal exposure (we use 30 quantiles). Lastly, for each quantile, we calculate the average observed loss cost, the average current pure premium, and the average new model predicted loss cost. In addition, we numerically examine predictive accuracy using a variety of validation measures. These measures are listed as follows: Gini index, percentage error (PE), root mean squared error (RMSE), and MAE. For definitions and interpretations of these validation measures, see Appendix C.

Take Figure 8(a) and Figure 8(b), the double lift charts of the BG group, for example. The observed loss cost (actual value) is indicated by the solid black line, the predictions provided by insurance in-house models are indicated by the red dotted line, and the predictions derived from InsurTech-enhanced LightGBM models are indicated by the blue dashed line. We can observe that the blue dashed line follows the trend of the solid black line closer, compared to the red dotted line, indicating that InsurTech-enhanced LightGBM models outperformed insurance in-house models with greatly improved prediction accuracy. Also, it can be seen that InsurTech-enhanced LightGBM models are able to capture the spike of the observed loss cost on the right side, which represents the larger claims. Similarly, from Figure 8(e) and Figure 8(f), we can see that the above statement also holds for the LIAB group. For the BP group, the InsurTech-enhanced LightGBM model exhibits a slight edge compared to the insurance in-house model. These empirical results show that the data created by InsurTech innovations can significantly improve the insurance company’s loss models and help differentiate the larger claims. Though the double lift charts offer a broad perspective on the predictive performance of the model, we explore their exact prediction accuracy in the

⁶This work utilizes resources supported by the National Science Foundation’s Major Research Instrumentation program, Grant 1725729, as well as the University of Illinois Urbana–Champaign. Retrieved from <https://wiki.ncsa.illinois.edu/display/ISL20/HAL+cluster>

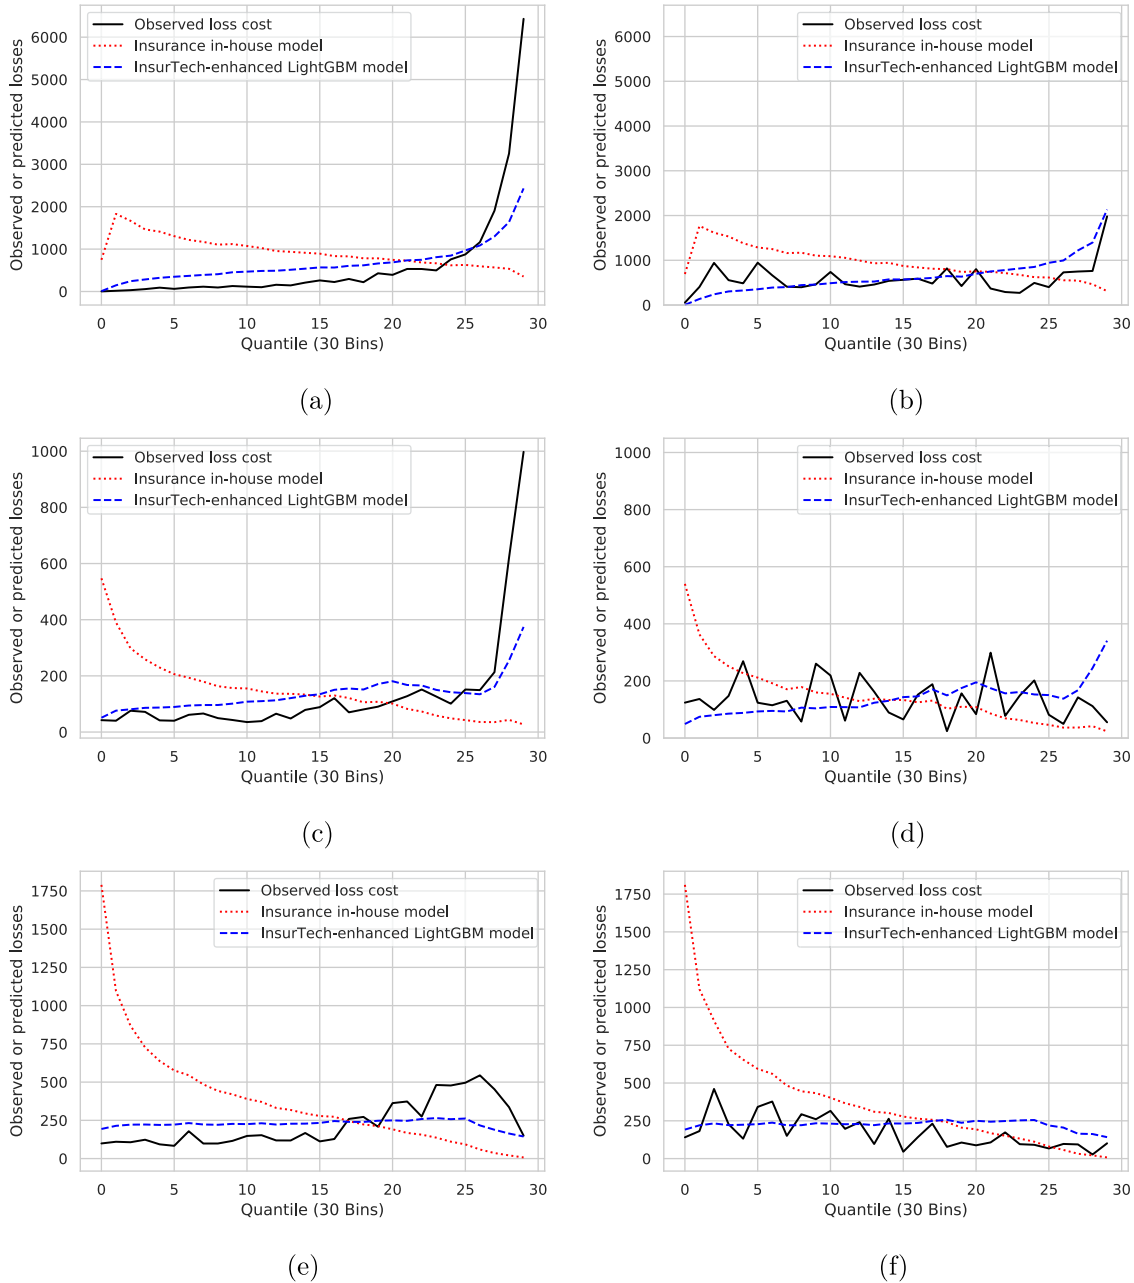


FIGURE 8. Double Lift Charts for Model Comparison: (a) Double Lift Chart for BG Train, (b) Double Lift Chart for BG Test, (c) Double Lift Chart for BP Train, (d) Double Lift Chart for BP Test, (e) Double Lift Chart for LIAB Train, and (f) Double Lift Chart for LIAB Test.

subsequent discussion by utilizing various validation metrics to gain a more profound understanding of the quantitative improvement of the model.

Table 4 compares the predictive performance of InsurTech-enhanced models, Tweedie GLM with elastic net feature selection and LightGBM, and the insurance in-house model based on the training and test datasets of each coverage group. For each coverage, we highlight the best-performing model with the value of each validation measure in bold. In general, both InsurTech-enhanced models consistently outperform the insurance in-house model for each coverage group. This suggests that the improvement stems from the additional information provided by InsurTech, irrespective of the chosen loss model. It should be noted that InsurTech-enhanced models significantly reduce the absolute value of PE of observed loss cost predictions compared to insurance in-house models. This indicates that InsurTech-enhanced models have superior predictive performance at

TABLE 4
Model Performance Based on Validation Measures

Coverage	Dataset	Model	Gini	PE	RMSE	MAE
BG	Train	Insurance in-house model	0.29	-0.40	5761.94	1526.47
		Tweedie GLM + elastic net	0.44	-0.04	5660.01	1286.31
		LightGBM	0.84	0.00	5364.05	1198.07
	Test	Insurance in-house model	0.32	-0.54	5328.02	1461.92
		Tweedie GLM + elastic net	0.32	-0.16	5284.90	1238.94
		LightGBM	0.37	-0.08	5198.57	1181.47
BP	Train	Insurance in-house model	0.59	-0.07	2498.13	277.37
		Tweedie GLM + elastic net	0.68	0.00	2450.82	262.64
		LightGBM	0.78	0.00	2409.88	259.11
	Test	Insurance in-house model	0.58	-0.11	2350.80	270.75
		Tweedie GLM + elastic net	0.36	-0.04	2375.10	262.31
		LightGBM	0.59	-0.06	2348.93	262.78
LIAB	Train	Insurance in-house model	0.57	-0.67	3937.22	586.88
		Tweedie GLM + elastic net	0.63	-0.04	3920.13	449.25
		LightGBM	0.78	0.00	3853.67	435.65
	Test	Insurance in-house model	0.54	-1.15	3347.60	547.02
		Tweedie GLM + elastic net	0.47	-0.33	3340.86	408.15
		LightGBM	0.56	-0.26	3305.85	394.56

the portfolio level, which is usually a key concern for insurers, particularly from a financial statement perspective. Concerning the comparable evaluation metrics observed in Tweedie GLM and LightGBM models, please refer to [Appendix C](#) for further discussion on the Rashomon effect.

4. MODEL INTERPRETATION AND DISCUSSION

Modern machine learning models have often been perceived as black-box models and are criticized for producing results that may be difficult to interpret, particularly in the actuarial domain. In this section, we demonstrate the actuarial interpretability of our calibrated models by examining feature importance and exploring a selection of illustrative business cases. Although this approach is not without its flaws, we still observe some insightful interpretations that resonate with our expert knowledge. Regarding our naming protocol, features with names preceded by uppercase letters, apart from the initial letter, are drawn from the InsurTech company files; all others have names that only have the initial letter capitalized and are drawn from the insurance company files.

4.1. Feature Importance

With more than 500 features in the dataset, we now evaluate how much each feature contributed to the prediction of the InsurTech-enhanced LightGBM model discussed in [Section 3](#). Three different approaches are applied to assess the importance of features: Mean Decrease in Impurity (MDI), Mean Decrease in Accuracy (MDA), and SHapley Additive exPlanations (SHAP). The MDI, also referred to as Gini feature importance, was detailed in Breiman et al. (1984). It measures the feature importance by calculating the average decrease in impurities of all splits in the tree-based model building process. The MDA, also called permutation feature importance, was introduced by Breiman (2001), which is a model-agnostic global explanation method that measures the importance of a feature by calculating the decrease in the model score (in our case, the increase in the model's mean squared error) after randomly permuting the given feature with other features. Lundberg and Lee (2017) proposed SHAP, which measures the local feature importance for every observation by computing Shapley values from coalitional game theory. SHAP is an additive feature attribution method whose model consists of a linear function of simplified binary indication of the entire feature space. The SHAP values are calculated by permuting the existence of each feature and the corresponding model predictions.




FIGURE 9. Top 20 Influential Features in InsurTech-Enhanced LightGBM Model (BG test): (a) Mean Decrease in Impurity—Top 20 Features, (b) Mean Decrease in Accuracy—Top 20 Features, and (c) SHAP Feature Importance—Top 20 Features.

Figure 9 highlights the top 20 features calculated by the three methods mentioned above in the InsurTech-enhanced LightGBM model on the BG subdataset. It illustrates the extent of each feature’s contribution to the prediction of the model, highlighting the significant features that may reveal potential rating factors that were overlooked in insurance pricing. There is a notable overlap in the top features based on the three different feature importance methods. These top features include the coordinates of addresses (*ADDR.longtitude* and *ADDR.latitude*) in category *business information*, proximity scores (i.e., *PROX.enter.dis*, *PROX.traffic.dis*, and *PROX.traffic.den*), and territory risks (i.e., *TERRITORY.xx*), all of which are related to

the location of a business; see Table 3. These location-related risk factors are fairly important to the BG coverage group, because they cover losses related to the building of a business. The territory risks, particularly those in version 2 (e.g., *TERRITORY.j2*, regarding fire risk characteristics), contribute significantly to the model’s predictive capability. This is not surprising and is expected, given that territory risk features are engineered from business risk characteristics and thus contain more risk information and have higher fill rates. Other top InsurTech features in the model include visibility indexes (i.e., *INDEX.visibility1* and *INDEX.visibility2*), review scores (i.e., *REVIEW.xx* extracted from text data), business classification (i.e., *CLASS.xx*, the overall business classification assigned to a business), and corresponding segment proportions (i.e., *PROP.xx*, the proportion of total votes for a segment given to a business). Furthermore, the coverage limit (i.e., *CovgLimit*) and risk type (i.e., *RiskType_xx*) provided by the insurance company are important rating factors that describe policy information and measure exposure information. Specifically, we can see from Figure 9(b) and Figure 9(c) that the feature risk type condo/office (i.e., *RiskType_ConOff*) is of high importance, which is consistent with our observations in Section 2.1 that the BOP policies of condo/office risk type have the highest observed loss cost on average. Please refer to Appendix D for further feature importance discussion on the Tweedie GLM model.

In a classical linear model, although it is challenging to avoid pitfalls when multicollinearity and high-dimensional data exist, we are used to interpreting the model using linear coefficients. For black-box models, we have similar tools to observe how one feature has an impact on the response variable by keeping other features unchanged. Apley and Zhu (2020) introduced accumulated local effects (ALE), which is a global model-agnostic explanation method that evaluates the average impact of a feature on the predictions by integrating the averaged local effect (i.e., the difference in the predictions across the conditional distribution). Compared to the widely used partial dependence plots, ALE is more robust against correlated features and less computationally expensive. Furthermore, there are other techniques to achieve similar goals, such as individual conditional expectation plots to display one line per observation.

Figure 10 presents the ALE plots of proximity traffic scores. The y-axis shows the centered effect on the predictions, which corresponds to the change in the model predictions when the feature is given a particular value relative to the average prediction. Figure 10(a) shows the ALE plot of the traffic density score, the bar charts represent the size of each category, and the dashed line represents the effect of each value of the density score. With a density score of 1, there is a relatively positive impact on the prediction; however, as the density score increases, the relative effect decreases and becomes negative. Note that this is in line with the definition of the traffic density score, where high scores indicate low risk. Similarly, Figure 10(b) presents a downward trend in the effect on predictions as the traffic distance score increases. Furthermore, there is a sharp decrease in the effect of the distance score when it exceeds 3.0, which could serve as an indicator of risk for insurers. This gives us a clearer visualization of how proximity traffic scores impact the observed claim loss compared to the one-way analysis in Figure 6. Specifically, in a one-way analysis, the effect of a particular feature on the response variable is evaluated without considering the effect of other features contained in the dataset. The ALE plot can solve this issue by using model results that take into account correlations in features.

4.2. Illustrative Individual Cases

To examine how the InsurTech risk factors affect the loss model locally, four real businesses are observed from a microscopic point of view. Because it helps enhance model interpretability, we use the SHAP value to explain every individual prediction based on the InsurTech-enhanced LightGBM model. SHAP values quantify the feature importance of all features and explain how each feature impacts individual predictions. Due to space constraints, we only highlight 20 important features among more than 500 features in the context of absolute SHAP values. In Figure 11, the four real businesses are extracted from our dataset to illustrate how the features influence the observed claim loss. Note that the features labeled red positively

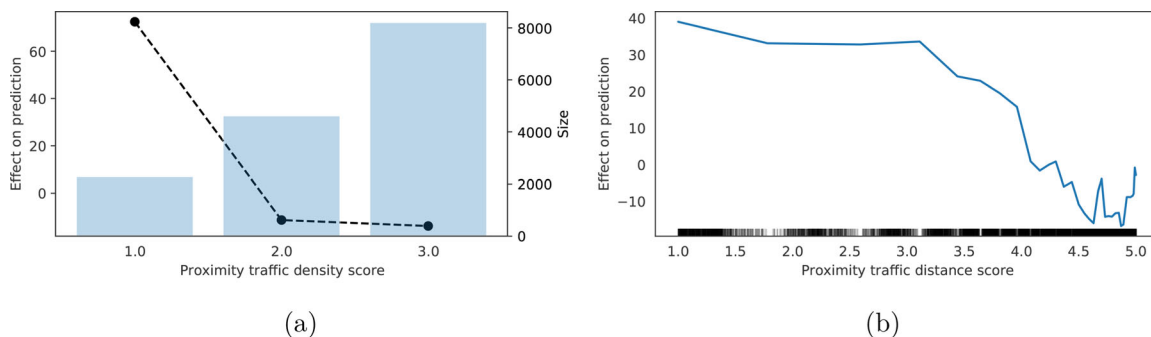


FIGURE 10. ALE Plots of Proximity Traffic Scores: (a) Proximity Traffic Density Score and (b) Proximity Traffic Distance Score.

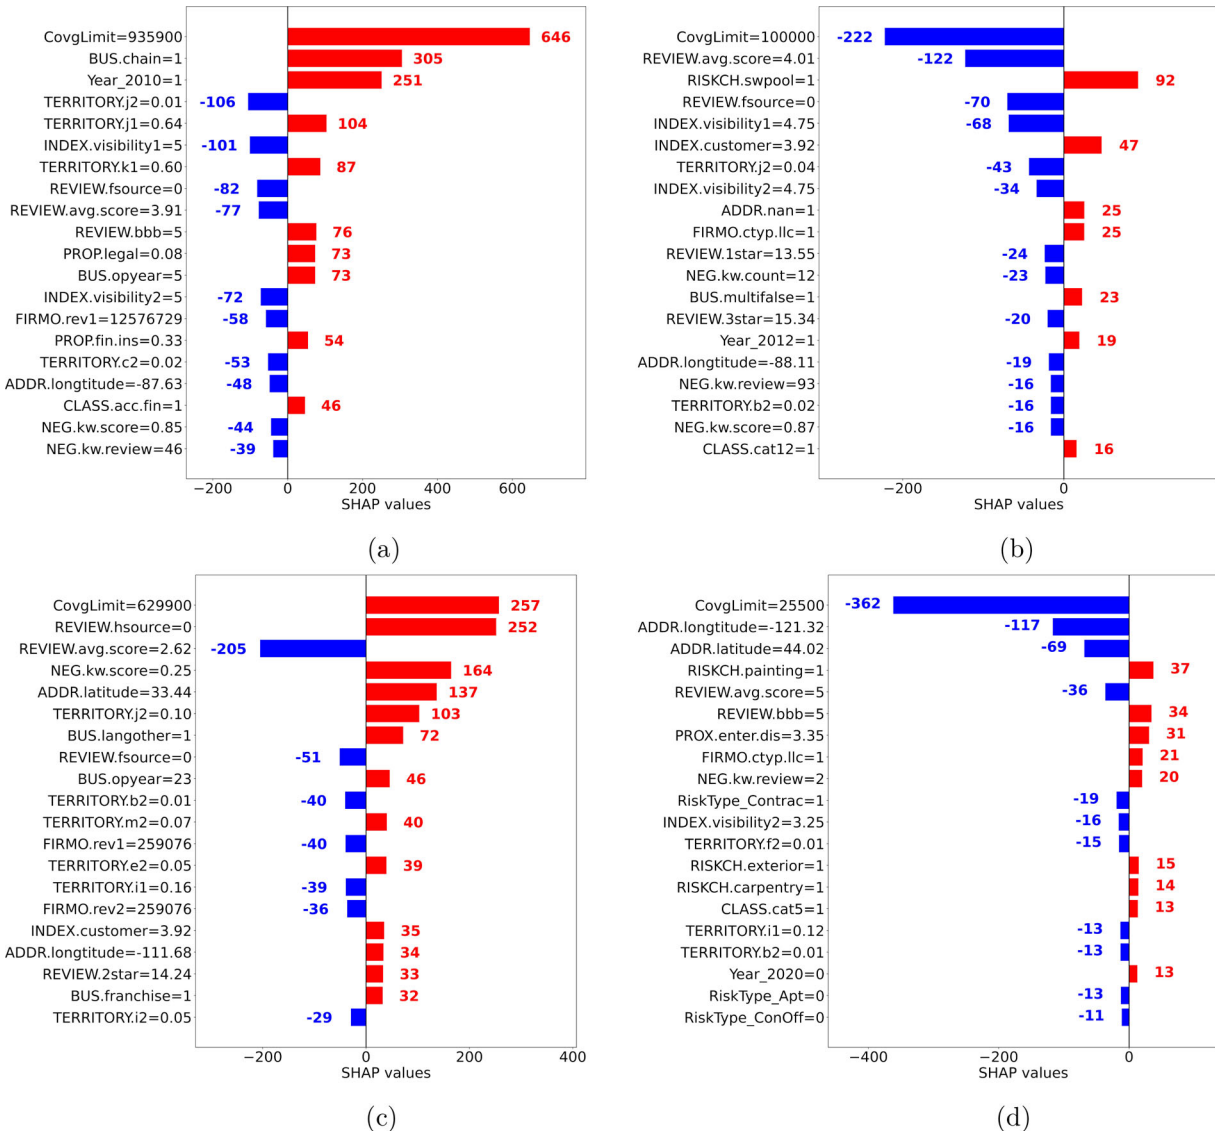


FIGURE 11. Top 20 Influential Features of four Local Examples: (a) Business A, (b) Business B, (c) Business C, and (d) Business D.

contribute to the observed claim loss, meaning this feature is relatively riskier than the features labeled blue. These businesses are described as follows (a) a business with a positive claim from the training dataset, (b) a business with no claim from the training dataset, (c) a business with a positive claim from the test dataset, (d) a business with no claim from the test dataset.

In (a), Business A is a land trust company that provides legal and financial services to real estate or property owners. InsurTech risk factor captures the corresponding business classification and segment information through *CLASS.acc.fin* (business classification is accountants & financial services), *PROP.legal* (business segment has the proportion of legal services), and *PROP.fin.ins* (the proportion of the finance & insurance), which belongs to the category *classification* in Table 3. Business A has multiple chain businesses statewide, which is indicated by the InsurTech risk factor *BUS.chain* in category *business information*, which may have a strong positive impact on the observed loss cost. In 2010, *Year_2010*, Business A purchased BOP from our insurance partner, after 5 years of operation, *BUS.opyear* = 5, which is considered a relatively young company with potential risk. Business A has a corporation revenue of \$12,576,729, based on *FIRMO.rev1* in category *firmographics*, which means that it has a more significant revenue than the average, thus being considered as a less risky business. Index and territory risk further describe the InsurTech risk factor associated with Business A, and these risk factors have a significant impact on determining the observed claim loss from a modeling perspective. We can also see that the *ADDR.longitude* shows up in the important feature list, because it is correlated with the territory risk, which is defined by ZIP code information. Similarly, social media and online content provide more detailed

information and its contribution to the response variable. For example, Business A has an average review score of 3.91, *REVIEW.avg.score*, the Better Business Bureau rating is 5, *REVIEW.bbb*, and these reviews exclude source F, *REVIEW.fsource*, and all the information comes from the category *text data*. InsurTech company also creates negative keyword scores and counts, *NEG.kw.score* and *NEG.kw.review*, to summarize these text data using natural language processing techniques.

Take another example, (b), Business B is a rental apartment. This business joined as a BOP policyholder in year 2012, implied by *YEAR_2012*, and has not submitted any claims during the period for which the policy is in force. As advertised on its website, the apartment premises include a swimming pool, a risk captured by *RISKCH.swpool* (business risk characteristics whether the business has a swimming pool) of category *risk characteristics*. Industries such as rental apartments are heavily influenced by subjective customer reviews and ratings. Engineered InsurTech features that quantify those scores can help identify the risk. In this example, the overall appearance of Business B across social media, ratings, and review sites entails to what extent the business values its targeted internet advertisement, which is measured by the engineered scores *INDEX.visibility1* and *INDEX.visibility2* of category *index*. Business B, with scores for both of 4.75, is higher than average and is considered a safer business in this regard. Similarly, another feature of category *index*, the customer ratings, *INDEX.customer*, is rated 3.92, a relative score, and is lower than expected from the perspective of a rental apartment, which explains the positive impact on claim assessment. Additionally, detailed score distributions and likes/dislikes (and comments) can increase the accuracy of risk identification. In the common 5-stars rating system where 1 star represents the lowest quality of service/products and 5 stars suggests the highest, the proportion of each of the star ratings can help distinguish the service provided, and that information is captured by InsurTech features *REVIEW.1star* through *REVIEW.5star* of the broader category *text data*. A 13.55 score of *REVIEW.1star*, which is lower than average, and a 15.34 score of *REVIEW.3star*, which is higher than expected, are consistent with a high average (4.01) score of *REVIEW.ave.score* summarized from previous customers' review contents for Business B. Overall, it has a safer business operation.

The claim submitted by Business C demonstrated in example (c) further strengthens the point that the insurance industry can benefit from InsurTech features for insurance risk assessment. Business C is a licensed medical clinic specializing in internal medicine that purchased a BOP policy to insure its business operations against potential risks. The insured clinic is a part of a franchise business, indicated by *BUS.franchise*, suggesting probable operation risk. The service business provided is also highly correlated with risk evaluation. As one of those identifying criteria, *BUS.langothor* is a feature that distinguishes whether the business supports less commonly used languages in the United States. In the example of Business C, this implies another complication in medical practice and thus a more risky business operation. Similar to previous examples, internet content plays a significant role for the business. Today, prospective patients have a habit of relying on internet reviews to help them decide which physicians to visit. The score, *REVIEW.hsource* of class *text data*, denotes the number of scores from Source H, which is an online information platform specialized for physicians and hospitals. Counted 0 from Business C indicates that fewer customer reviews appear than expected and, correspondingly, a significant positive impact on claim prediction.

Belonging to one of the major business lines, contractors, Business D in (d) is a painting company that, as advertised, provides interior/exterior painting and color design services. The business, as its legal status stands, is a Limited Liability Company (LLC), suggested by *FIRMO.ctyp.llc*, which differentiates whether the company is an LLC. The limited liability property suggests potential risk, which contributes positively to loss prediction. The features *RISKCH.painting*, *RISKCH.exterior*, and *RISKCH.carpentry* from the same category *risk characteristics* denote the identification of risk characteristics such as whether the business provides painting, exterior design, and carpentry services, respectively, and accurately capture the characteristics of the business. The score in category *proximity score*, *PROX.enter.dis*, is an engineered feature quantifying nearby entertainment-related risk measured by distance. Business D has a higher risk related to entertainment than expected, whereas other location-based risk is lower and therefore considered a safe business.

Across all four examples, some of the other InsurTech features reveal a consistent impact on predictions. Both Businesses B and D are categorized as LLCs, and both contribute positively to loss predictions. Visibility scores based on features *INDEX.visibility1*, *INDEX.visibility2* and customer ratings *INDEX.customer* can be found in all four examples, and a higher overall appearance of the business on the internet (higher visibility score) or higher customer feedback can be interpreted as lower business operation risk and thus a negative influence on final predictions. In the common 5-star rating systems where 5 stars stands for the highest quality of products/services represented by the features *REVIEW.1star* through *REVIEW.5star*, as seen from Businesses B and C, a higher proportion of low stars (1 star or 2 stars) implies less trust from customers and thus a higher possibility of operation risk. Conversely, a higher proportion of positive reviews suggests a safer business operation.

In addition to InsurTech features, features such as coverage limit *CovgLimit* and groups/risk types provided by insurance companies and used as rating factors still play a significant role in each of the interpretations, because they are comparatively important in describing policy content. However, to better characterize individual policies, InsurTech features outperform many of the traditional insurance features that the insurance company uses as rating factors in our empirical study.

5. CONCLUDING REMARKS

The continuing growth in InsurTech innovations is transforming the insurance industry in its manner of embracing rapid advances in technology to modernize the entire insurance value chain. Such modernization results in the flow of big data that allows InsurTech to offer innovative data-driven solutions across all major lines of business in the insurance industry. This article illustrates how to leverage InsurTech innovations to enhance the forecasting of business insurance loss costs and quantifies the improvements from such innovations. Using a portfolio comprising nearly a million BOP policies, this work not only proposes advancements in InsurTech-enhanced loss modeling but also advocates for the integrating benefits derived from both traditional rating factors used by an insurance company and external risk factors collected by an InsurTech company. A critical challenge is combining and centralizing these data sources, and through a successful academic–industry collaboration of the University of Illinois Urbana–Champaign, Carpe Data (an InsurTech company), and an anonymous insurance company, such work is made possible.

An InsurTech-enhanced loss model plays a pivotal role in an insurance company’s quest to reduce the loss ratio, and advancements in this model, which this article proposes, significantly help a company manage risks more effectively. Its capacity to accurately assess risk enables insurers to evaluate the associated risks of diverse policies and policyholders with greater precision. Such precision extends to premium determination, ensuring that policyholders are charged amounts closely aligned with expected losses, thereby minimizing the risk of underpricing. Early identification of potential claims allows insurers to proactively mitigate risks, implement preventive measures, or adjust premiums. Furthermore, the adaptability of an InsurTech-enhanced loss model to changing market conditions and emerging risks enables timely adjustments to underwriting and pricing strategies, which can be crucial for staying ahead of evolving risks, preventing adverse selection, and maintaining a balanced portfolio for a lower loss ratio over time. In summary, an enhanced loss model improves an insurance company’s ability to assess risk accurately, set appropriate premiums, and proactively manage potential losses, ultimately leading to a more favorable loss ratio essential for financial health and profitability.

We envision that insurance companies may face challenges when adopting the work proposed here to seek regulatory approval for premium rating and risk classification. Given the surge of artificial intelligence and big data, it is imperative to acknowledge the evolving landscape of insurance. Regulators must devise a framework to address various issues, including the tendency for large insurance companies to reap greater benefits from InsurTech, potentially becoming more resilient to adverse selection. As highlighted in Brunnermeier, Lamba, and Segura-Rodriguez (2021), insurers may leverage new data sources or technologies to gain informational advantages over policyholders, potentially leading to concerns about price discrimination and inverse selection. However, such challenges can be readily overcome with the proper justification. Model interpretability revealed that many relevant and intuitive features that affect loss costs are derived from the InsurTech dataset, including firmographics, risk characteristics, indexes, proximity, and territory risk scores. In addition, using visualization, feature importance, and illustrative individual cases may help improve the understanding of modern machine learning algorithms. From a regulatory standpoint, policymakers may consider mandating insurers to disclose all features they possess, thereby rendering private and advantageous information accessible through a public data repository. Conversely, consumer activism and raising policyholders’ awareness about data utilization could prove more effective in curtailing exploitation through inverse selection than mandating data to be public. Additionally, fostering competition serves as a means to safeguard policyholders, mitigating the extent of price discrimination by big data monopolies.


ACKNOWLEDGMENT

The authors thank the anonymous insurance company for sharing the BOP insurance data, Carpe Data for providing InsurTech-empowered risk information, and the National Center for Supercomputing Applications for offering computing resources and technical support. In particular, the authors thank Dan, Derek, Eli, Eric, Lois, Michael, Mingxi, Perri, Shuijing, Steven, Shirui Luo (NCSA), and Volodymyr Kindratenko (NCSA) for their expertise and assistance. Please note that the last names of the above contributors are omitted due to privacy issues. We further thank the Editor of the *North American Actuarial Journal*, Yijia Lin, and the anonymous referees for providing constructive input that helped us significantly improve the article.

DISCLOSURE STATEMENT

No potential conflict of interest was reported by the author(s).

ORCID

Zhiyu Quan  <http://orcid.org/0000-0001-9735-6713>
 Changyue Hu  <http://orcid.org/0000-0003-0975-4012>
 Panyi Dong  <http://orcid.org/0000-0003-0957-0141>
 Emiliano A. Valdez  <http://orcid.org/0000-0003-4037-8385>

REFERENCES

- Akiba, T., S. Sano, T. Yanase, T. Ohta, and M. Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD'19*, 2623–31. New York, NY: Association for Computing Machinery.
- Apley, D. W., and J. Zhu. 2020. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82 (4):1059–86. doi:10.1111/rssb.12377
- Berman, T., D. Schallmo, and C. A. Williams. 2021. Business model transformation through artificial intelligence in the Israeli InsurTech. Paper presented at ISPIIM Connects Valencia, Valencia, Spain, November 29–December 1.
- Blier-Wong, C., L. Lamontagne, and E. Marceau. 2024. A representation-learning approach for insurance pricing with images. *ASTIN Bulletin* 54 (2):280–309. doi:10.1017/asb.2024.9
- Boucher, J.-P., S. Côté, and M. Guillen. 2017. Exposure as duration and distance in telematics motor insurance using generalized additive models. *Risks* 5 (4):54. doi:10.3390/risks5040054
- Bowsher, C. G., and P. S. Swain. 2012. Identifying sources of variation and the flow of information in biochemical networks. *Proceedings of the National Academy of Sciences* 109 (20):E1320–28. doi:10.1073/pnas.1119407109
- Breiman, L. 2001. Random forests. *Machine Learning* 45 (1):5–32. doi:10.1023/A:1010933404324
- Breiman, L., J. Friedman, C. Stone, and R. Olshen. 1984. *Classification and regression trees*. Monterey, CA: Wadsworth & Brooks/Cole Advanced Books & Software.
- Brunnermeier, M. K., R. Lamba, and C. Segura-Rodriguez. 2021. Inverse selection. doi:<http://dx.doi.org/10.2139/ssrn.3584331>
- Che, X., A. Liebenberg, and J. Xu. 2022. Usage-based insurance—Impact on insurers and potential implications for InsurTech. *North American Actuarial Journal* 26 (3):428–55. doi:10.1080/10920277.2021.1953536
- Chen, Z., Y. Lu, J. Zhang, and W. Zhu. 2023. Managing weather risk with a neural network-based index insurance. *Management Science* 70 (7):4306–27. doi:10.1287/mnsc.2023.4902
- Chester, A., N. Hoffmann, S. Johansson, and P. B. Olesen. 2018. Commercial lines InsurTech: A pathway to digital. Technical report, McKinsey & Company, New York, NY, USA.
- Delgado-Panadero, Á., B. Hernández-Lorca, M. T. García-Ordás, and J. A. Benítez-Andrades. 2022. Implementing local-explainability in gradient boosting trees: Feature contribution. *Information Sciences* 589:199–212. doi:10.1016/j.ins.2021.12.111
- Denuit, M., A. Charpentier, and J. Trufin. 2021. Autocalibration and Tweedie-dominance for insurance pricing with machine learning. *Insurance: Mathematics and Economics* 101:485–97. doi:10.1016/j.insmatheco.2021.09.001
- Denuit, M., X. Maréchal, S. Pitrebois, and J.-F. Walhin. 2007. *Actuarial modelling of claim counts: Risk classification, credibility and bonus-malus systems*. Chichester, UK: John Wiley & Sons.
- Diao, L., and C. Weng. 2019. Regression tree credibility model. *North American Actuarial Journal* 23 (2):169–96. doi:10.1080/10920277.2018.1554497
- Federal Insurance Office, U.S. Department of the Treasury (2023). Annual Report on the Insurance Industry. Accessed 2023. <https://home.treasury.gov/system/files/311/FIO%20Annual%20Report%202023%209292023.pdf>.
- Financial Technology Partners. 2023. Insurtech insights reports. Accessed 2023. <https://www.ftpartners.com/fintech-research/insurtech-almanac>.
- Fluckiger, I., and M. Duygun. 2022. New technologies and data in insurance. *The Geneva Papers on Risk and Insurance - Issues and Practice* 47:495–98. doi:10.1057/s41288-022-00274-6
- Frees, E. W. J., and P. Shi. 2017. Credibility prediction using collateral information. *Variance* 11 (1):45–59.
- Friedman, J. H. 2001. Greedy function approximation: A gradient boosting machine. *Annals of Statistics* 29 (5):1189–232.
- Gao, G., S. Meng, and M. V. Wüthrich. 2019. Claims frequency modeling using telematics car driving data. *Scandinavian Actuarial Journal* 2019 (2):143–62. doi:10.1080/03461238.2018.1523068
- Guelman, L. 2012. Gradient boosting trees for auto insurance loss cost modeling and prediction. *Expert Systems with Applications* 39 (3):3659–67. doi:10.1016/j.eswa.2011.09.058
- Guillen, M., J. P. Nielsen, M. Ayuso, and A. M. Pérez-Marín. 2019. The use of telematics devices to improve automobile insurance rates. *Risk Analysis* 39 (3):662–72. doi:10.1111/risa.13172
- Guillen, M., J. P. Nielsen, and A. M. Pérez-Marín. 2021. Near-miss telematics in motor insurance. *Journal of Risk & Insurance* 88 (3):569–89. doi:10.1111/jori.12340
- Guillen, M., J. P. Nielsen, A. M. Pérez-Marín, and V. Elpidorou. 2020. Can automobile insurance telematics predict the risk of near-miss events? *North American Actuarial Journal* 24 (1):141–52. doi:10.1080/10920277.2019.1627221
- Hastie, T., R. Tibshirani, and J. H. Friedman. 2009. *The elements of statistical learning: Data mining, inference, and prediction*, vol. 2. New York, NY: Springer.
- Henckaerts, R., M.-P. Côté, K. Antonio, and R. Verbelen. 2021. Boosting insights in insurance tariff plans with tree-based machine learning methods. *North American Actuarial Journal* 25 (2):255–85. doi:10.1080/10920277.2020.1745656
- Hu, C., Z. Quan, and W. F. Chong. 2022. Imbalanced learning for insurance using modified loss functions in tree-based models. *Insurance: Mathematics and Economics* 106:13–32. doi:10.1016/j.insmatheco.2022.04.010
- Huang, Y., and S. Meng. 2019. Automobile insurance classification ratemaking based on telematics driving data. *Decision Support Systems* 127:113156. doi:10.1016/j.dss.2019.113156

- Jørgensen, B., and M. C. Paes De Souza. 1994. Fitting Tweedie's compound Poisson model to insurance claims data. *Scandinavian Actuarial Journal* 1994 (1):69–93. doi:10.1080/03461238.1994.10413930
- Ke, G., Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu. 2017. LightGBM: A highly efficient gradient boosting decision tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 3149–57. Red Hook, NY: Curran Associates.
- Kelley, C., and K. Wang. 2021. Insurtech: A guide for the actuarial community. Technical report, Willis Tower Watson. Published by the Society of Actuaries.
- Kindratenko, V., D. Mu, Y. Zhan, J. Maloney, S. H. Hashemi, B. Rabe, K. Xu, R. Campbell, J. Peng, and W. Gropp. 2020. Hal: Computer system for scalable deep learning. In *Practice and Experience in Advanced Research Computing 2020: Catch the Wave*, PEARC '20, 41–48. New York, NY: Association for Computing Machinery. 10.1145/3311790.3396649
- Lee, S. C., and S. Lin. 2018. Delta boosting machine with application to general insurance. *North American Actuarial Journal* 22 (3):405–25. doi:10.1080/10920277.2018.1431131
- Loh, W.-Y. 2011. Classification and regression trees. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 1 (1):14–23. doi:10.1002/widm.8
- Lundberg, S. M., and S.-I. Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, eds. U. V. Luxburg, I. Guyon, S. Bengio, H. Wallach, and R. Fergus, 4765–74. Red Hook, NY: Curran Associates.
- McCrea, M., and M. Farrell. 2018. A conceptual model for pricing health and life insurance using wearable technology: Pricing insurance using wearable technology. *Risk Management and Insurance Review* 21 (3):389–411. doi:10.1111/rmir.12112
- Naylor, M. 2017. Types of insurance. In *Insurance transformed: Technological disruption*, 41–45. Cham, Switzerland: Palgrave Macmillan.
- Nelder, J. A., and R. W. Wedderburn. 1972. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)* 135 (3):370–84. doi:10.2307/2344614
- Organisation for Economic Co-operation and Development. 2023. Annual GDP and components - Expenditure approach. Accessed 2023. https://stats.oecd.org/Index.aspx?DatasetCode=SNA_TABLE1.
- Paes, L. M., R. Cruz, F. P. Calmon, and M. Diaz. 2023. On the Inevitability of the Rashomon Effect. In *2023 IEEE International Symposium on Information Theory (ISIT)*, 549–554.
- Pesantez-Narvaez, J., M. Guillen, and M. Alcañiz. 2019. Predicting motor insurance claims using telematics data—XGBoost versus logistic regression. *Risks* 7 (2):70. doi:10.3390/risks7020070
- Qian, W., Y. Yang, and H. Zou. 2016. Tweedie's compound Poisson model with grouped elastic net. *Journal of Computational and Graphical Statistics* 25 (2):606–25. doi:10.1080/10618600.2015.1005213
- Qiao, Y., C.-W. Wang, and W. Zhu. 2024. Machine learning in long-term mortality forecasting. *The Geneva Papers on Risk and Insurance - Issues and Practice* 49:340–62. doi:10.1057/s41288-024-00320-5
- Quan, Z., and E. A. Valdez. 2018. Predictive analytics of insurance claims using multivariate decision trees. *Dependence Modeling* 6 (1):377–407. doi:10.1515/demo-2018-0022
- Ricci, O., and F. Battaglia. 2021. The development of InsurTech in Europe and the strategic response of incumbents. In *Disruptive technology in banking and finance: An international perspective on FinTech*, eds. T. King, F. S. Stentella Lopes, A. Srivastav, and J. Williams, 135–60. Cham, Switzerland: Springer International Publishing.
- Shahriari, B., K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE* 104 (1):148–75. doi:10.1109/JPROC.2015.2494218
- Song, Y.-y., and Y. Lu. 2015. Decision tree methods: Applications for classification and prediction. *Shanghai Archives of Psychiatry* 27 (2):130–35. doi:10.11919/j.issn.1002-0829.215044
- Spender, A., C. Bullen, L. Altmann-Richer, J. Cripps, R. Duffy, C. Falkous, M. Farrell, T. Horn, J. Wigzell, and W. Yeap. 2019. Wearables and the Internet of Things: Considerations for the life and health insurance industry. *British Actuarial Journal* 24:1–31. doi:10.1017/S1357321719000072
- Suryavanshi, U. 2022. The InsurTech revolution in insurance industry: Emerging trends, challenges and opportunities. *International Journal of Management and Development Studies* 11:12–19. doi:10.53983/ijmds.v11n08.002
- The Donegal Insurance Group. (2022). Donegal annual reports. Accessed 2023. <https://investors.donegalgroup.com/financial-reporting/annual-reports-proxies>.
- The Hanover Insurance Group. (2022). Hanover annual reports. Accessed 2023. <https://investors.hanover.com/financials/annual-reports/default.aspx>.
- The Travelers Companies. (2022). Travelers annual reports. Accessed 2023. <https://investor.travelers.com/financial-information/annual-reports/default.aspx>.
- VanderLinden, S. L., S. M. Millie, N. Anderson, and S. Chishti. 2018. *The INSURTECH book: The insurance technology handbook for investors, entrepreneurs and FinTech visionaries*. Chichester, UK: John Wiley & Sons.
- Verbelen, R., K. Antonio, and G. Claeskens. 2018. Unravelling the predictive power of telematics data in car insurance pricing. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 67 (5):1275–304.
- Wang, Q. 2021. The impact of InsurTech on Chinese insurance industry. *Procedia Computer Science* 187:30–35. doi:10.1016/j.procs.2021.04.030
- Weidner, W., F. W. G. Transchel, and R. Weidner. 2017. Telematic driving profile classification in car insurance pricing. *Annals of Actuarial Science* 11 (2):213–36. doi:10.1017/S1748499516000130
- Wuthrich, M. V., and C. Buser. 2021. Data analytics for non-life insurance pricing. Swiss Finance Institute Research Paper Series 16–68, Swiss Finance Institute, Zürich, Switzerland.
- Xu, X., and P. Zweifel. 2020. A framework for the evaluation of InsurTech. *Risk Management and Insurance Review* 23 (4):305–29. doi:10.1111/rmir.12161
- Yang, Y., W. Qian, and H. Zou. 2018. Insurance premium prediction via gradient tree-boosted Tweedie compound Poisson models. *Journal of Business & Economic Statistics* 36 (3):456–70. doi:10.1080/07350015.2016.1200981
- Zhou, H., W. Qian, and Y. Yang. 2022. Tweedie gradient boosting for extremely unbalanced zero-inflated data. *Communications in Statistics - Simulation and Computation* 51 (9):5507–29. doi:10.1080/03610918.2020.1772302
- Zhu, W. 2024. A deep factor model for crop yield forecasting and insurance ratemaking. *North American Actuarial Journal* 28 (1):57–72. doi:10.1080/10920277.2023.2182792
- Zou, H., and T. Hastie. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 67:301–20. doi:10.1111/j.1467-9868.2005.00503.x

APPENDIX A. TRADITIONAL LOSS MODELS

A.1. Tweedie Generalized Linear Model

Generalized linear models (GLMs), proposed by Nelder and Wedderburn (1972), introduce a linear or non-linear relationship between features and response variables through a link function and a distribution function. GLM is a widely accepted actuarial practice because of its advantages in implementation and interpretation, both of which are motivated by the fact that insurance is a strictly regulated industry. The traditional practice of actuarial ratemaking is to find proper loss distributions under the GLM framework. Because the distribution of insurance loss typically contains a point mass at zero representing no claims and a continuous component (often substantial) for positive values representing the amount of a claim, Tweedie GLM has been one of the most frequently used actuarial ratemaking models in the past few decades. See Jørgensen and Paes De Souza (1994).

With a general loss model framework,

$$Y = \sum_{i=1}^N Z_i, \quad (\text{A.1})$$

where N refers to the number of claims and Y denotes the total claim amount, so that Z_i captures the claim severity of each claim. The Tweedie distribution, which is categorized as a class of compound Poisson-Gamma distribution, is a special case of such a general loss model where N is Poisson distributed and Z_i ($i = 1, 2, \dots, N$) are independently identically distributed Gamma random variables. The Tweedie distribution is a subclass of the family of exponential dispersion models, which can be written as

$$f(y, \mu, \tau, \phi) = a(y, \phi) \exp \left(\frac{1}{\phi} \left(\frac{y\mu^{1-\tau}}{1-\tau} - \frac{\mu^{2-\tau}}{2-\tau} \right) \right), \quad (\text{A.2})$$

where ϕ is the dispersion parameter in $(0, +\infty)$, mean $E(Y) = \mu$ and variance $\text{Var}(Y) = \phi\mu^\tau$, τ is an extra parameter that controls the variance of the distribution. In the special case of $p = 0$, the distribution degenerates to the Normal distribution; when $p = 1$, it becomes a Poisson distribution; when $p = 2$, it is equivalent to Gamma distribution; when $p = 3$, it is an inverse Gaussian distribution.

Here, we assume that the dataset consists of a vector of p features, x_1, x_2, \dots, x_p , and a response variable y ; that is, our data can be represented as

$$(\mathbf{X}, Y) = ((\mathbf{X}^{IH}, \mathbf{X}^{IT}), Y) = \{\mathbf{x}_i, y_i\}_i = \{(x_{i1}, x_{i2}, \dots, x_{ip}), y_i\}_i, \quad (\text{A.3})$$

where $i = 1, \dots, n$ and n is the number of observations. Once the choice of distribution and link function (logarithmic link) is identified, the regression coefficients, β_i ($i = 1, 2, \dots, p$), can be found by the maximum likelihood estimation. Specifically, under the Tweedie GLM framework, the negative log-likelihood can be written as

$$l(\boldsymbol{\beta}) = \sum_{i=1}^n w_i \left(\frac{y_i e^{-(\tau-1) \sum_{i=1}^p x_i \beta_i}}{\tau-1} + \frac{e^{(2-\tau) \sum_{i=1}^p x_i \beta_i}}{2-\tau} \right), \quad (\text{A.4})$$

where w_i is the weight for each observation (by default, equal weight is assigned); here, it assumes that ϕ is the same for all observations, which is questionable in our real-life dataset. Furthermore, we use our dataset to estimate the mean by $\log(\mu) = \sum_{i=1}^p x_i \beta_i$.

However, given the characteristics of our real-life dataset, the Tweedie GLM is not an ideal model, because we have over 500 risk factors requiring an overwhelming feature selection task, and there is also a problem of missing data, which is almost inevitable, demanding extensive imputation tasks.

A.2. Tweedie Generalized Linear Model with Elastic Net

Qian, Yang, and Zou (2016) examined the group elastic net regularization for Tweedie GLM and proposed an efficient algorithm to model high-dimensional datasets. In our experiments, we also utilized elastic net regularization to reduce the dimension of the feature space. Initially proposed by Zou and Hastie (2005), the elastic net has served as a variable

selection method through a combination of L1 and L2 regularization on regression coefficients. Under the same setting as in [Appendix A.1](#), the elastic net minimizes the objective function

$$\boldsymbol{\beta} = \operatorname{argmin}_{\boldsymbol{\beta}} l(\boldsymbol{\beta}) + \lambda \left[(1 - \alpha) \frac{1}{2} \|\boldsymbol{\beta}\|_{\ell_2}^2 + \alpha \|\boldsymbol{\beta}\|_{\ell_1} \right], \quad (\text{A.5})$$

where $l(\boldsymbol{\beta})$ is defined as [Equation \(A.4\)](#), with L1 norm $\|\boldsymbol{\beta}\|_{\ell_1}$ and L2 norm $\|\boldsymbol{\beta}\|_{\ell_2}$ of the regression coefficients, and $\lambda \geq 0$ and $0 \leq \alpha \leq 1$ are tuning parameters. If $\alpha = 1$, the optimization reduces to lasso regression and if $\alpha = 0$, the optimization reduces to ridge regression.

After standardizing all of the features, we assume that the more influential features have larger coefficients in absolute values. To achieve the goal of feature selection, variables whose coefficients generated by elastic net larger than some pre-defined threshold β^{thres} are selected for further modeling. After feature selection, we use the subset of our dataset to build the Tweedie GLM. This modeling strategy has two benefits: first, we preserve the original feature scale when we build the Tweedie GLM, and the coefficients are more interpretable compared to directly using elastic net modeling; second, the Tweedie GLM is one of the current ratemaking tools approved by regulators.

A.3. Decision Trees: Classification and Regression Tree

We follow the notation used in [Hastie et al. \(2009\)](#). The base learner, DT, is an algorithm that partitions data into M regions through the feature space recurrently by layers of binary internal nodes into disjoint leaf nodes R_1, \dots, R_M , where each leaf node R_m is assigned a constant c_m ($m = 1, 2, \dots, M$) as prediction. This indicates that a dataset (\mathbf{X}, Y) can be modeled by

$$f(\mathbf{x}_i) = \sum_{m=1}^M c_m \mathbb{1}(\mathbf{x}_i \in R_m), \quad (\text{A.6})$$

where $\mathbb{1}$ is an indicator function and $\mathbb{1}(\mathbf{x}_i \in R_m)$ is equal to 1 when $\mathbf{x}_i \in R_m$. Each leaf node R_m consists of the conjunction of decision regions of the preceding internal nodes and the constant c_m of prediction that minimizes node impurity. The structure of internal nodes and leaf nodes is grown as the result of recursive binary splitting and stopping criterion. See [Loh \(2011\)](#).

Classification and regression tree ([Breiman et al. 1984](#)), in the case of regression, optimizes the sum of squared loss

$$\sum_{i:\mathbf{x}_i \in R_m} (y_i - f(\mathbf{x}_i))^2 \quad (\text{A.7})$$

as the impurity of node R_m , for which it solves for the estimation of constant c_m in the form

$$\hat{c}_m = \frac{1}{|R_m|} \sum_{i:\mathbf{x}_i \in R_m} y_i = \operatorname{ave}(y_i | \mathbf{x}_i \in R_m), \quad (\text{A.8})$$

where $|R_m|$ refers to the cardinality of the node R_m and $\operatorname{ave}(\cdot)$ denotes the average. The splitting procedure is designed to gain the most information possible throughout the split. The split at the initial node is to find such j ($j = 1, 2, \dots, p$) indicating feature x_j and decision boundary s so that in the disjoint split regions $R_{left} = \{\mathbf{x}_i | x_{ij} \leq s\}$ and $R_{right} = \{\mathbf{x}_i | x_{ij} > s\}$, the objective function, controlled by j and s , is minimized. The objective function can be expressed as

$$\operatorname{arg\,min}_{j,s} \sum_{i:\mathbf{x}_i \in R_{left}} (y_i - \hat{c}_{R_{left}})^2 + \sum_{i:\mathbf{x}_i \in R_{right}} (y_i - \hat{c}_{R_{right}})^2, \quad (\text{A.9})$$

where $\hat{c}_{R_{left}} = \frac{1}{|R_{left}|} \sum_{i:\mathbf{x}_i \in R_{left}} y_i$, $\hat{c}_{R_{right}} = \frac{1}{|R_{right}|} \sum_{i:\mathbf{x}_i \in R_{right}} y_i$, respectively. For the following internal nodes, the same splitting procedure is applied recursively to the subregions split by preceding decision rules to find the optimal feature and the decision boundary within the subregions. The tree is then grown after each successful split and gained information represented by the global loss reduction. As indicated by [Song and Lu \(2015\)](#), to prevent an overly complex tree structure and

overfitting, and to maintain time efficiency, stopping rules like the minimum number of observations in a leaf node and the depth of a tree are practical ones to regularize individual tree sizes.

APPENDIX B. MODEL HYPERPARAMETER SETTING

Explanation of hyperparameters optimized in the Tweedie GLM after elastic net feature selection

1. *alpha*: *alpha* is the constant that controls penalization by multiplying both the L1 norm and the L2 norm regularization term.
2. *l1_ratio*: *l1_ratio* is proportion of L1 norm penalization. Correspondingly, the L2 norm is regularized by $1-l1_ratio$. If $l1_ratio = 0$, the elastic net becomes equivalent to the ridge regression; if $l1_ratio = 1$, the elastic net reduces to the lasso regression; if $0 < l1_ratio < 1$, it is a mixture of L1 and L2 regularization.
3. *coef_threshold*: *coef_threshold* is the threshold of feature selection. Variables whose absolute values of coefficients are larger than this threshold are selected for the Tweedie GLM.
4. *p*: *p* refers to the variance power that decides the variance of Tweedie distribution.

Explanation of hyperparameters optimized in LightGBM⁷:

1. *feature_fraction*: *feature_fraction* is a constant of range $0 < feature_fraction \leq 1$ referring to the proportion of features used for each tree growing. When $feature_fraction = 1$, all features are used.
2. *learning_rate*: *learning_rate* is a positive real number that controls the learning step of each tree update.
3. *max_depth*: *max_depth* denotes the maximum number of layers of each tree. This hyperparameter controls the size of trees to reduce training time and prevent overfitting.
4. *min_child_samples*: *min_child_samples* denotes the minimum number of samples needed for further node split. The hyperparameter is also critical to control the size of trees.
5. *num_leaves*: *num_leaves* indicates the maximum number of leaf nodes allowed in a tree. This hyperparameter limits the complexity of each tree in the ensemble.
6. *reg_alpha*: *reg_alpha* denotes the scale of L1 regularization.
7. *reg_lambda*: *reg_lambda* refers to the parameter of L2 penalization.
8. *subsample*: *subsample* is the hyperparameter referring to the bagging fraction where only a proportion of data is sampled for each tree split to speed up the training process.

TABLE B.1
Model Hyperparameter Setting

Model	Hyperparameter	Value (BG)	Value (BP)	Value (LIAB)
Tweedie GLM + elastic net	<i>alpha</i>	0.0249	0.1016	1.2602
	<i>l1_ratio</i>	0.0860	0.2228	0.3844
	<i>coef_threshold</i>	1.0	1.0	0.1
	<i>p</i>	1.1341	1.0004	1.3357
LightGBM	<i>feature_fraction</i>	0.5944	0.4482	0.4177
	<i>learning_rate</i>	0.0114	0.0455	0.0023
	<i>max_depth</i>	28	29	24
	<i>min_child_samples</i>	80	80	150
	<i>num_leaves</i>	120	40	50
	<i>reg_alpha</i>	0.4056	0.0892	0.7631
	<i>reg_lambda</i>	0.9800	0.9626	0.8913
	<i>subsample</i>	0.5192	0.5215	0.9319
	<i>subsample_for_bin</i>	120,000	100,000	60,000
	<i>n_estimators</i>	181	67	999

⁷Full reference of hyperparameters by LightGBM. Retrieved from <https://lightgbm.readthedocs.io/en/latest/Parameters.html/>

9. `subsample_for_bin`: `subsample_for_bin` is the number of samples used to construct discrete feature bins.
10. `n_estimators`: `n_estimators` implies the number of trees or iterations in the boosting ensemble. More trees grown represent more complex the boosting ensemble.

APPENDIX C. VALIDATION MEASURES AND RASHOMON EFFECT

- **Gini Index.** $\text{Gini} = 1 - \frac{2}{N-1} \left(N - \frac{\sum_{i=1}^N \tilde{y}_i}{\sum_{i=1}^N \hat{y}_i} \right)$, where \tilde{y} represents the corresponding observed value y based on the ranking of the predicted values \hat{y} . The higher the Gini, the better.
- **Percentage Error.** $\text{PE} = \frac{\sum_{i=1}^N \hat{y}_i - \sum_{i=1}^N y_i}{\sum_{i=1}^N y_i}$. The lower the PE, the better.
- **Root Mean Squared Error.** $\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2}$. The lower the RMSE, the better.
- **Mean Absolute Error.** $\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i|$. The lower the MAE, the better.

As observed from Table 4, LightGBM outperforms Tweedie GLM in most test data validation metrics, with the exception of PE and MAE in BP coverage. Nevertheless, both InsurTech-enhanced models perform comparably. This phenomenon potentially indicates the Rashomon effect, as suggested by Paes et al. (2023). The Rashomon effect occurs when multiple machine learning model architectures yield similar validation metrics, pointing to the presence of multiple, equally valid models that can explain the same data. Following on Paes et al. (2023), we can construct an upper bound for Rashomon bound for each coverage and confidence level $1 - \delta$,

$$|\hat{L}(h_1) - \hat{L}(h_2)| \leq 2\sqrt{\frac{\log 2/\delta}{2n}}, \quad (\text{A.10})$$

where h_1 and h_2 are the two models being compared, \hat{L} denotes the empirical loss associated with each input model, and n is the data size. Table C.1 demonstrates the constructed upper bounds. When the absolute difference of two empirical losses is lower than the constructed upper bound, the two models are considered δ -indistinguishable and possess comparable predictive power. If the absolute difference exceeds the upper bound, then there is a statistical difference in their predictive capabilities. By examining the Gini index, which is bounded by one, we can see that, at the 99.9% confidence level, there is a statistical difference in model performance, with LightGBM outperforming across all three coverages.

TABLE C.1
Rashomon Effect Bound

Coverage	Confidence level	Rashomon bound
BG	95%	0.0143
	99%	0.0171
	99.9%	0.0205
BP	95%	0.0108
	99%	0.0129
	99.9%	0.0155
LIAB	95%	0.0090
	99%	0.0108
	99.9%	0.0130

APPENDIX D. FURTHER DISCUSSION ON TWEEDIE GLM MODEL

As detailed in Section 4, the InsurTech enhancement significantly deepens our understanding of potential risk characteristics. To further illustrate this, we examine the feature importance of the Tweedie GLM model by extracting the top 20 absolute t -values, in the form

$$|t_j| = \left| \frac{\hat{\beta}_j}{\sigma_{\hat{\beta}_j}} \right|,$$

where $\hat{\beta}_j$ is the estimated GLM coefficient of the j th feature and $\sigma_{\hat{\beta}_j}$ is the standard deviation of the corresponding feature. See Figure D.1. The higher the absolute t -values, the more influential the feature is.

Similar to our findings in Section 4, territory risks (i.e., *TERRITORY.xx*) continue to be a significant presence in the top 20 influential features. Additionally, business classification (i.e., *CLASS.xx*) and review scores (i.e., *REVIEW.xx*) remain among the most influential features, impacting the Tweedie GLM model significantly. Insurance policy information such as policy year (i.e., *YEAR_xx*) and risk type (i.e., *RiskType_xx*) are also considered highly influential, helping to segment insurance policies accurately. The similarity in feature importance between LightGBM and Tweedie GLM models validates the significance of InsurTech-enhanced features. However, some differences, such as the absence of coverage limit (e.g., *CovgLimit*), the most significant feature in LightGBM, indicate that model output can be dependent on the chosen model structure. Therefore, selecting appropriate model architectures remains meaningful in practice.




FIGURE D.1. Top 20 Influential Features in InsurTech-Enhanced Tweedie GLM Model (BG).