Insurance in the Age of AI

A report series of the Geneva Association

Ruo (Alex) JIA

Deputy Editor-in-Chief, The Geneva Papers on Risk and Insurance – Issues and Practice

Associate Professor of Insurance, Peking University

2 July 2026

Seoul , 70th Industry-Academia Seminar at the Korea Insurance Research Institute

Geneva Association Gen AI and insurance report series

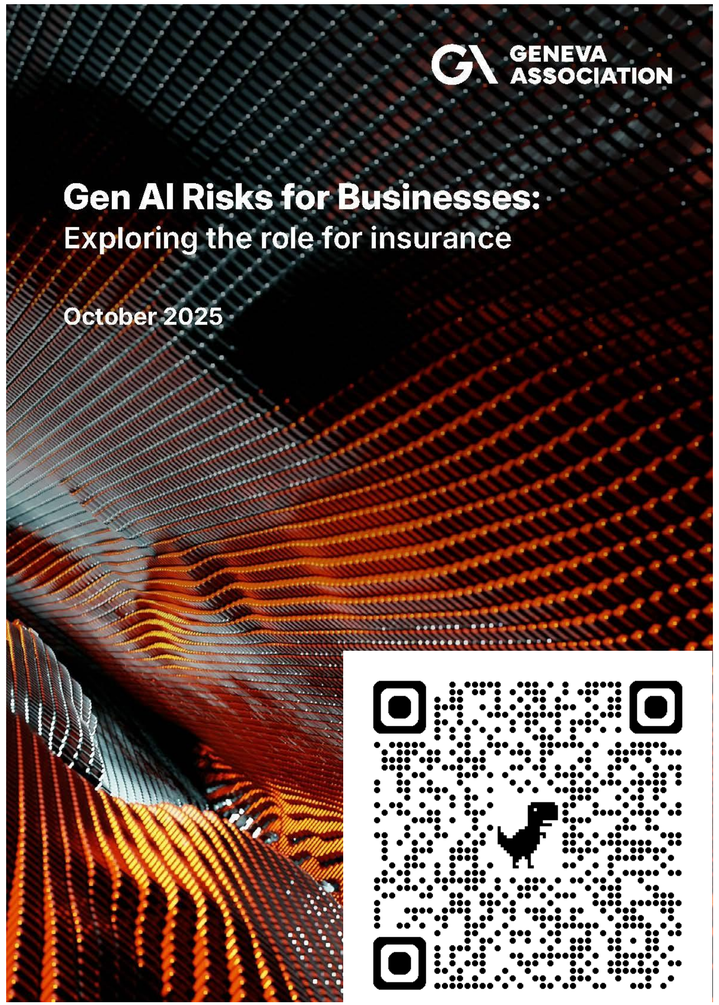

Report 1: [Gen AI Risks for Businesses: Exploring the role for insurance] (published on 2 October, 2025)

Report 2: [Gen AI in the Insurance Customer Journey]

(published on 20 November, 2025)

What is Generative AI – From Prediction to Creation

Traditional AI
The Analyst

- **Function:** Analyzes data to find patterns.
- **Output:** Predictions (e.g., “Is this transaction fraudulent?”).
- **Mechanism:** Deterministic mapping of predictors.
- **Insurability:** Clearer parameters, easier to validate.

Generative AI
The Creator

- **Function:** Creates NEW content (Text, Code, Audio, Video).
- **Output:** Original artifacts.
- **Mechanism:** Probabilistic; predicts next “token” in a sequence.
- **Insurability:** ‘Black Box’ opacity makes error tracing difficult.

Two Forms of Gen AI

Insurer-Provided Tools

Specialized tools deployed directly by insurance carriers.

Examples:

- Advanced chatbots
- Underwriting assistants
- Automated claims triage (e.g., Tokio Marine)
- Email drafting tools (e.g., Allstate)

Primary Goal:
Automation and personalized customer engagement.

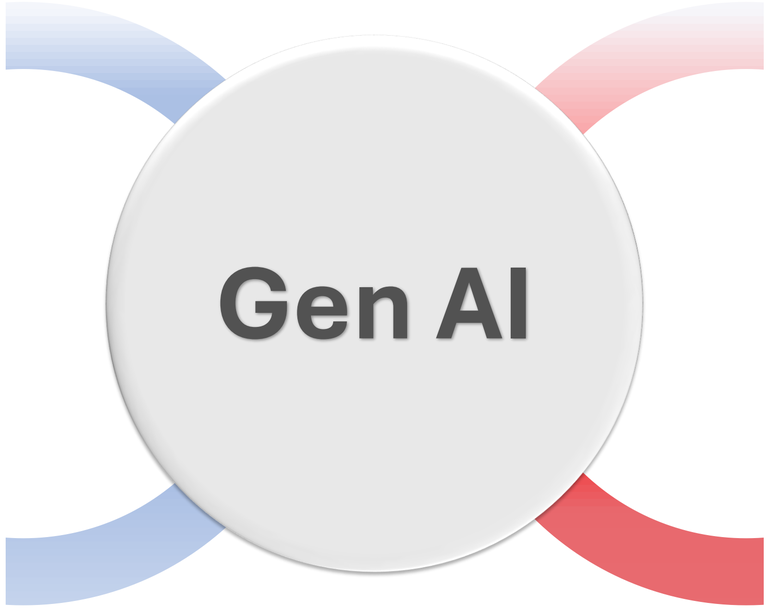

Off-the-Shelf Tools

General-purpose platforms (e.g., ChatGPT, Copilot, DeepSeek) used independently by customers.

Functions:

- Analyzing insurance policies
- Comparing products
- Clarifying terms without insurer involvement

Primary Goal:
Information gathering and decision support.

Impact Across the Customer Journey

Stages of the Insurance Journey:
- Information gathering
- Seeking advice
- Purchasing
- After-sales service
- Claims handling

Gen AI impact is strongest in:
- Advisory functions
- Claims communication
- After-sales engagement

	Gathering information	Seeking advice	Purchasing policy	After-sales service	Making a claim
Existing digital technologies (internet, mobile, social media, etc.)	High	Low → Medium	Medium	High	Medium
Additional Impact of Gen AI	High	High	Low → Medium	Medium → High	Medium → High

Customer Survey Overview

Data Source: Survey of 6,000 customers (Age 21–70).

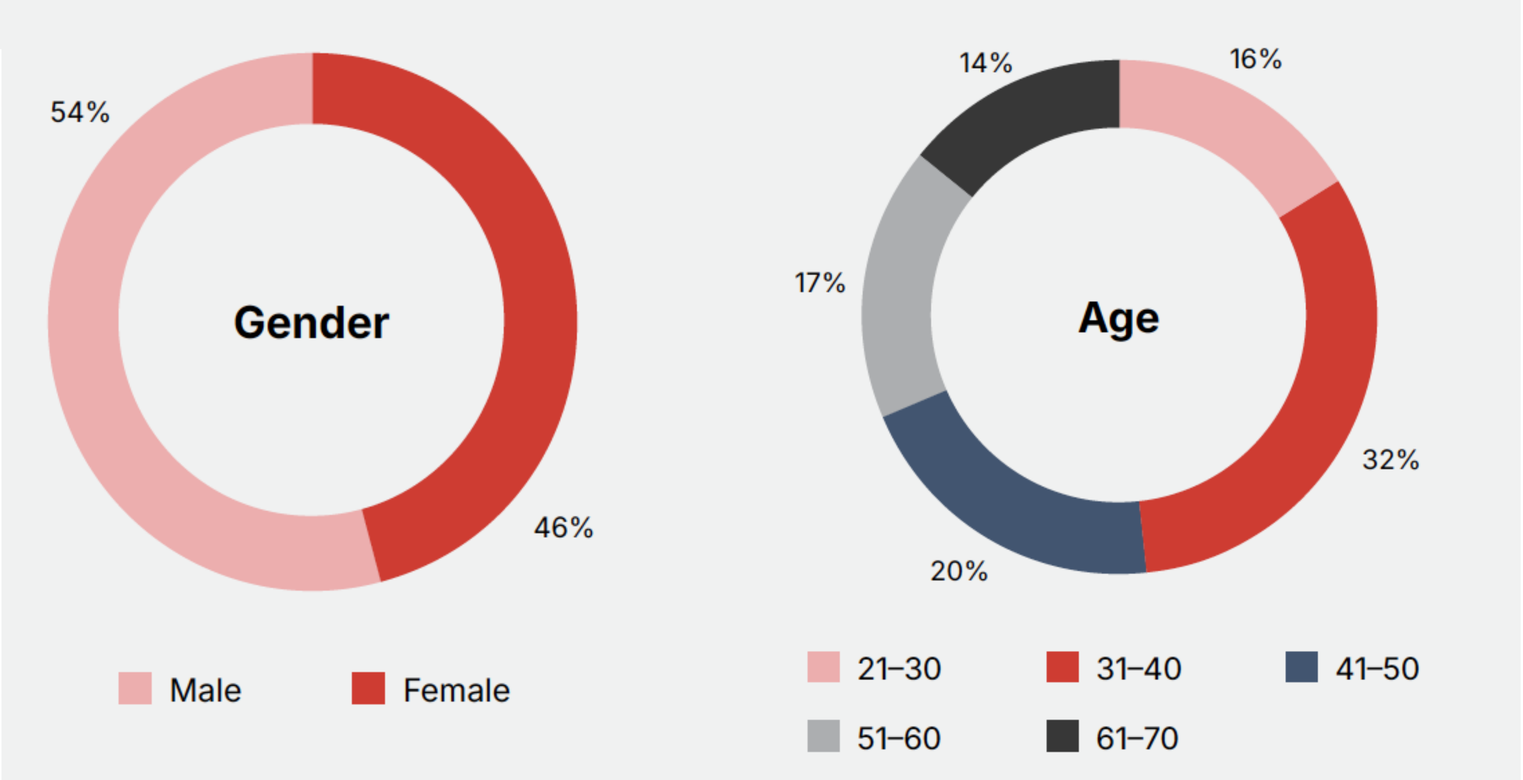

Must understand what Gen AI is.

Source: Geneva Association insurance customer survey

Customer Awareness of AI Use

- 45% noticed chatbots / virtual assistants
- 32% noticed auto-generated documents
- 18% unsure whether AI is used

Which, if any, of the following Gen AI applications have you noticed your insurance company using when interacting with you?

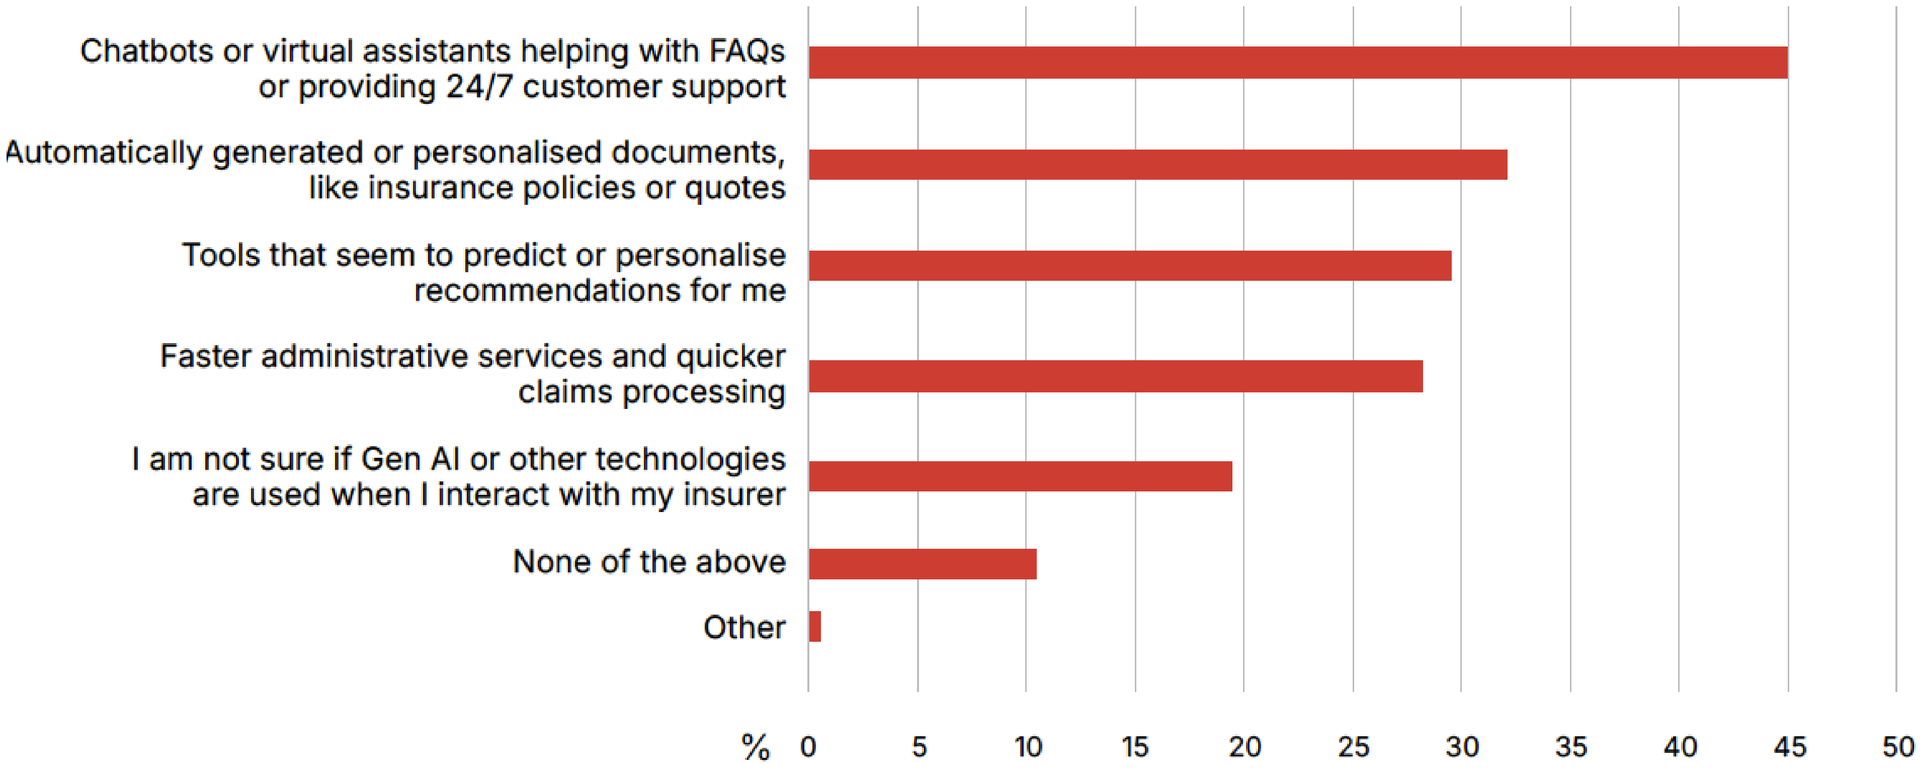

Source: Geneva Association insurance customer survey

Customer Sentiment Toward Insurer provided AI:

How do you feel about your insurer using Gen AI to interact with you?

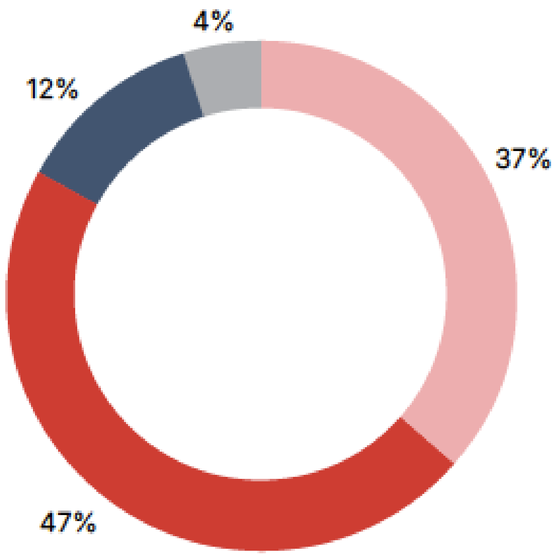

- I like it
- I am OK with it, as long as it provides services comparable to or better than human interactions
- I am not comfortable with it
- Not sure

Based on your experience, how would you describe the impact of Gen AI tools on your interactions with your insurer?

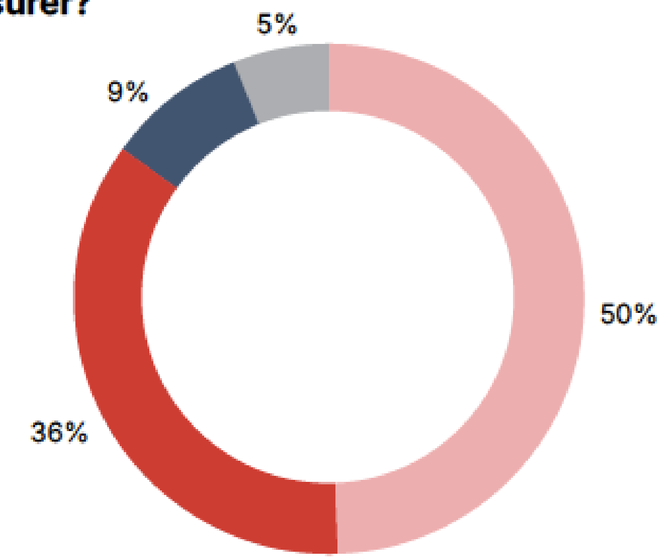

- They made my experience much easier and more efficient
- They were somewhat helpful but did not significantly improve my experience
- They made things more complicated or less personal
- I did not notice much of a difference

The East-West Divide in Acceptance

Asian markets exhibit significantly higher enthusiasm compared to Western skepticism.

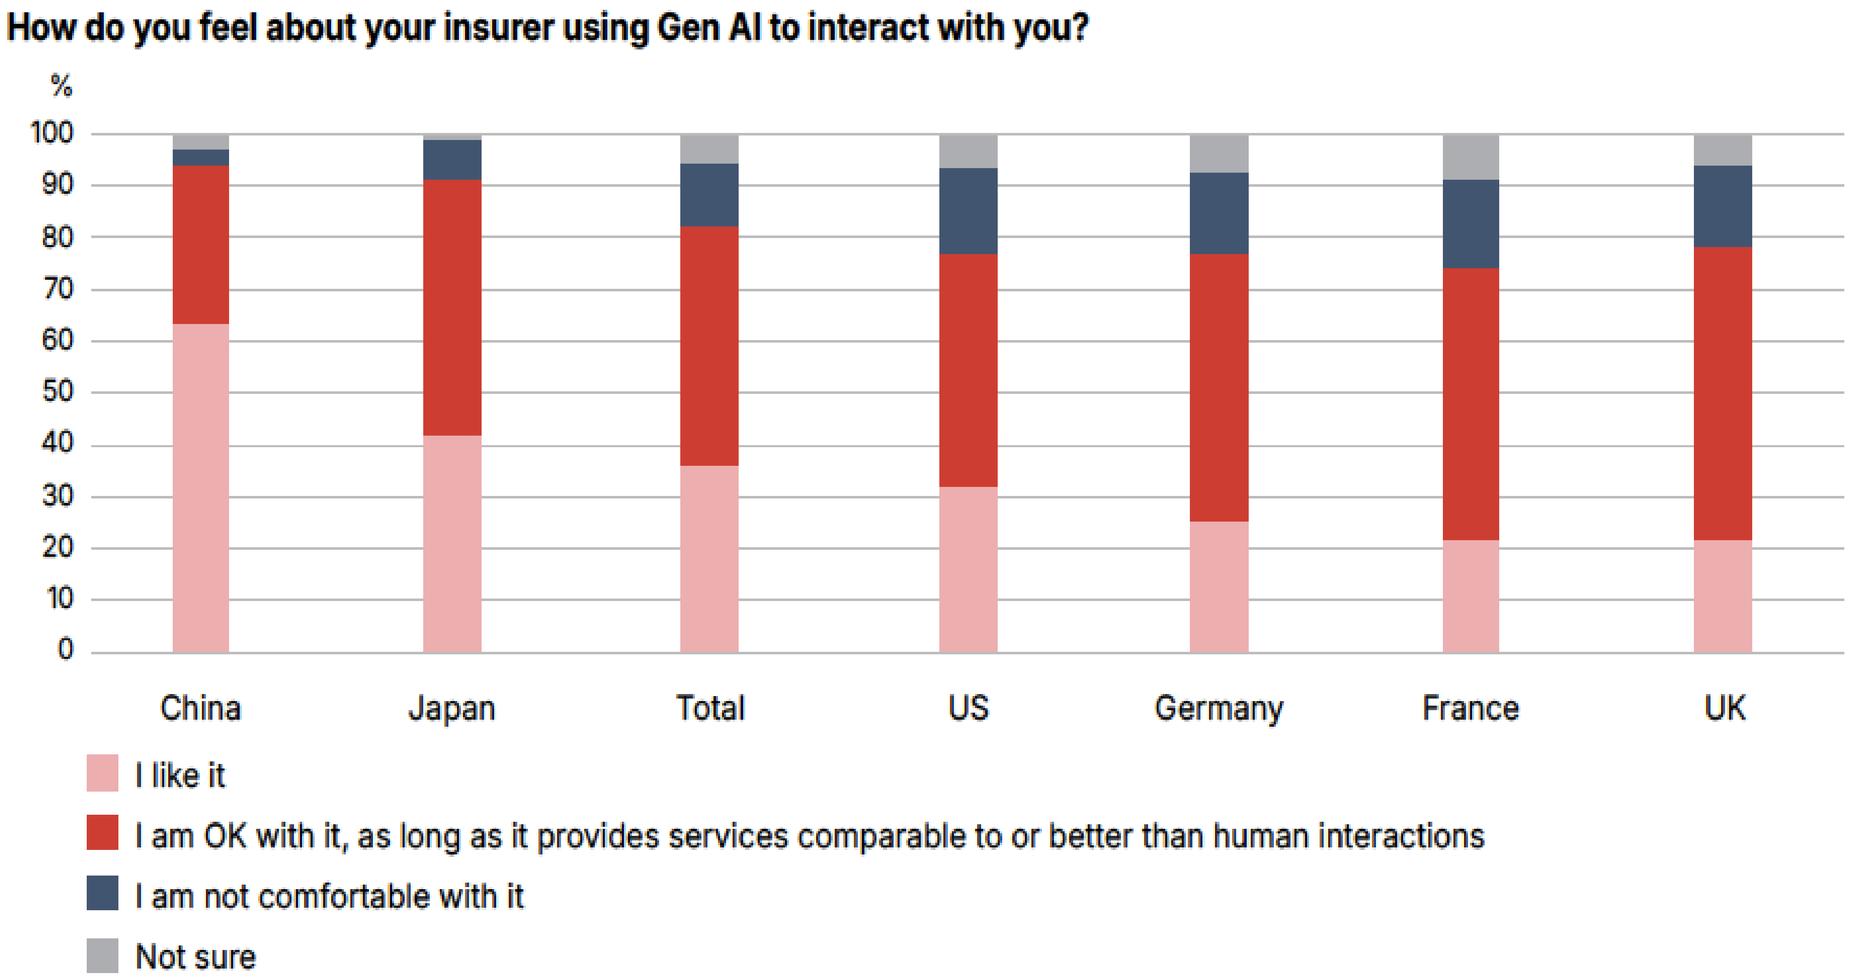

Regional Drivers

•**Asia**: High acceptance driven by digital ecosystem exposure and cultural openness.

Demand by Insurance Activity

AI is welcomed for utility, not advertisement.

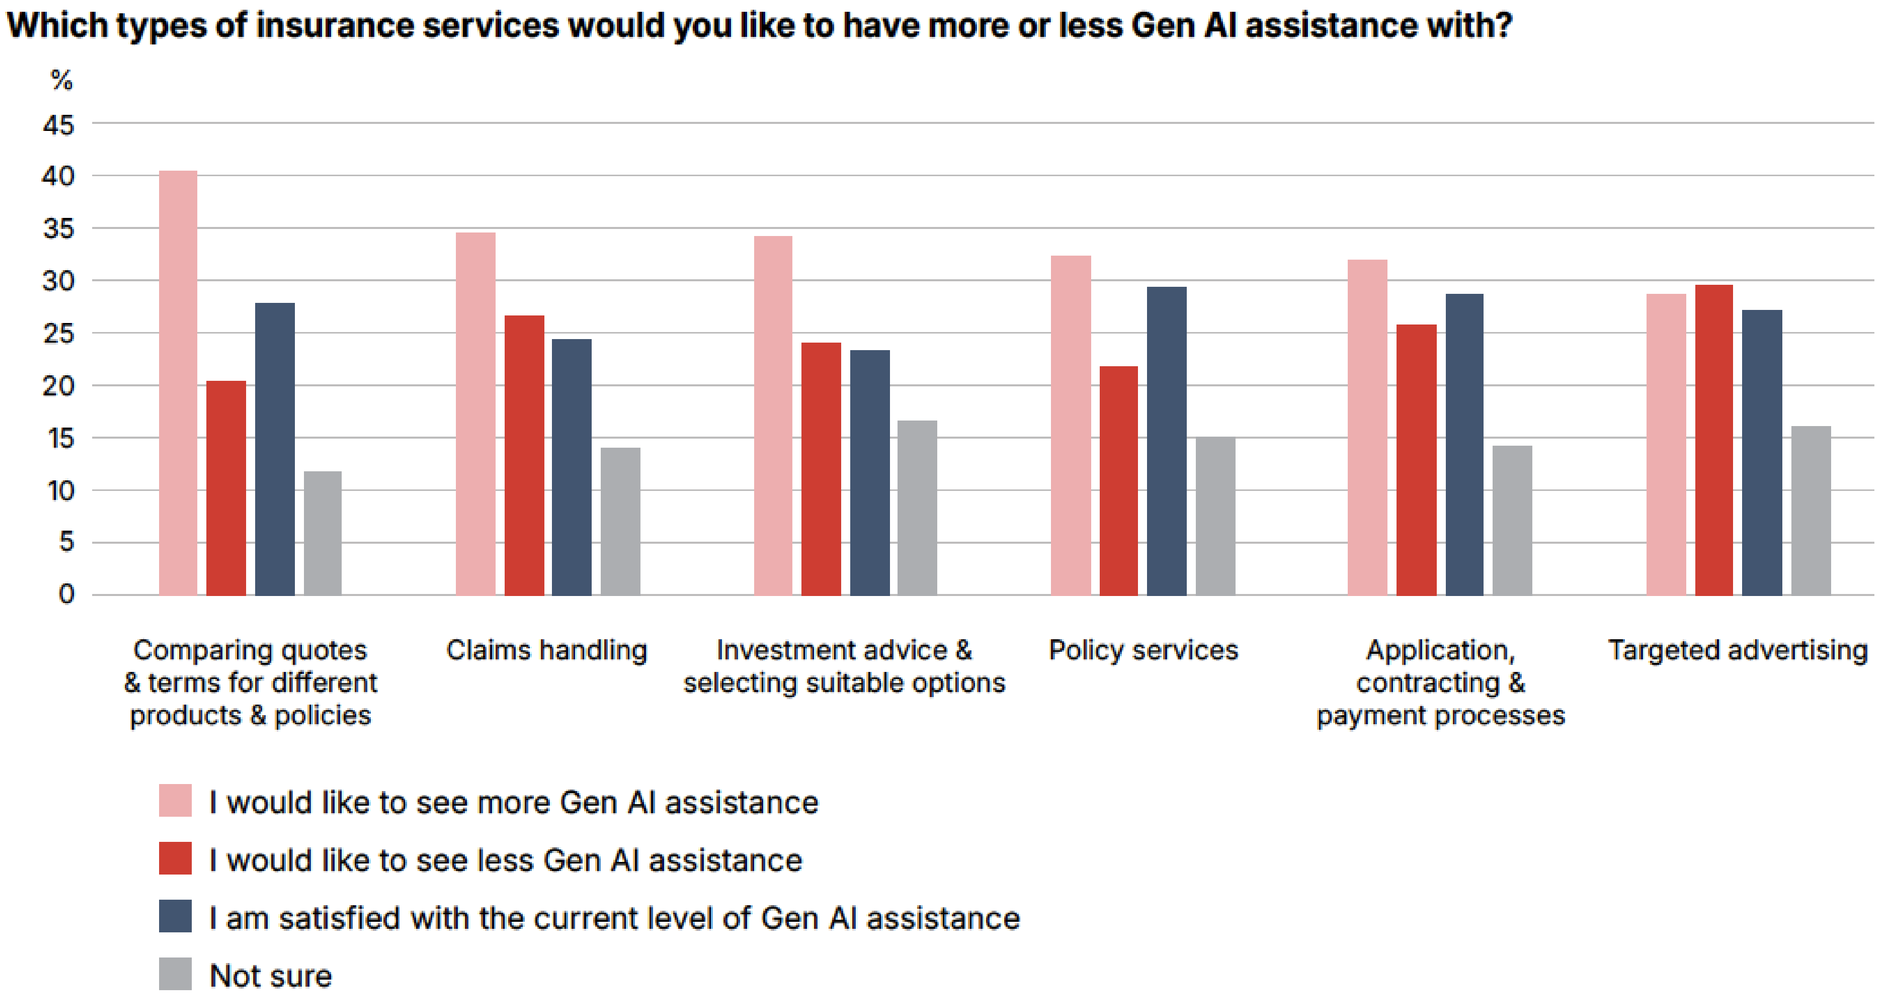

Customers want **MORE** AI in:

- Comparing quotes
- Claims handling
- Policy servicing

Customers want **LESS** AI in:

- Targeted advertising

Top Customer Priorities

Trust-related factors dominate speed or cost.

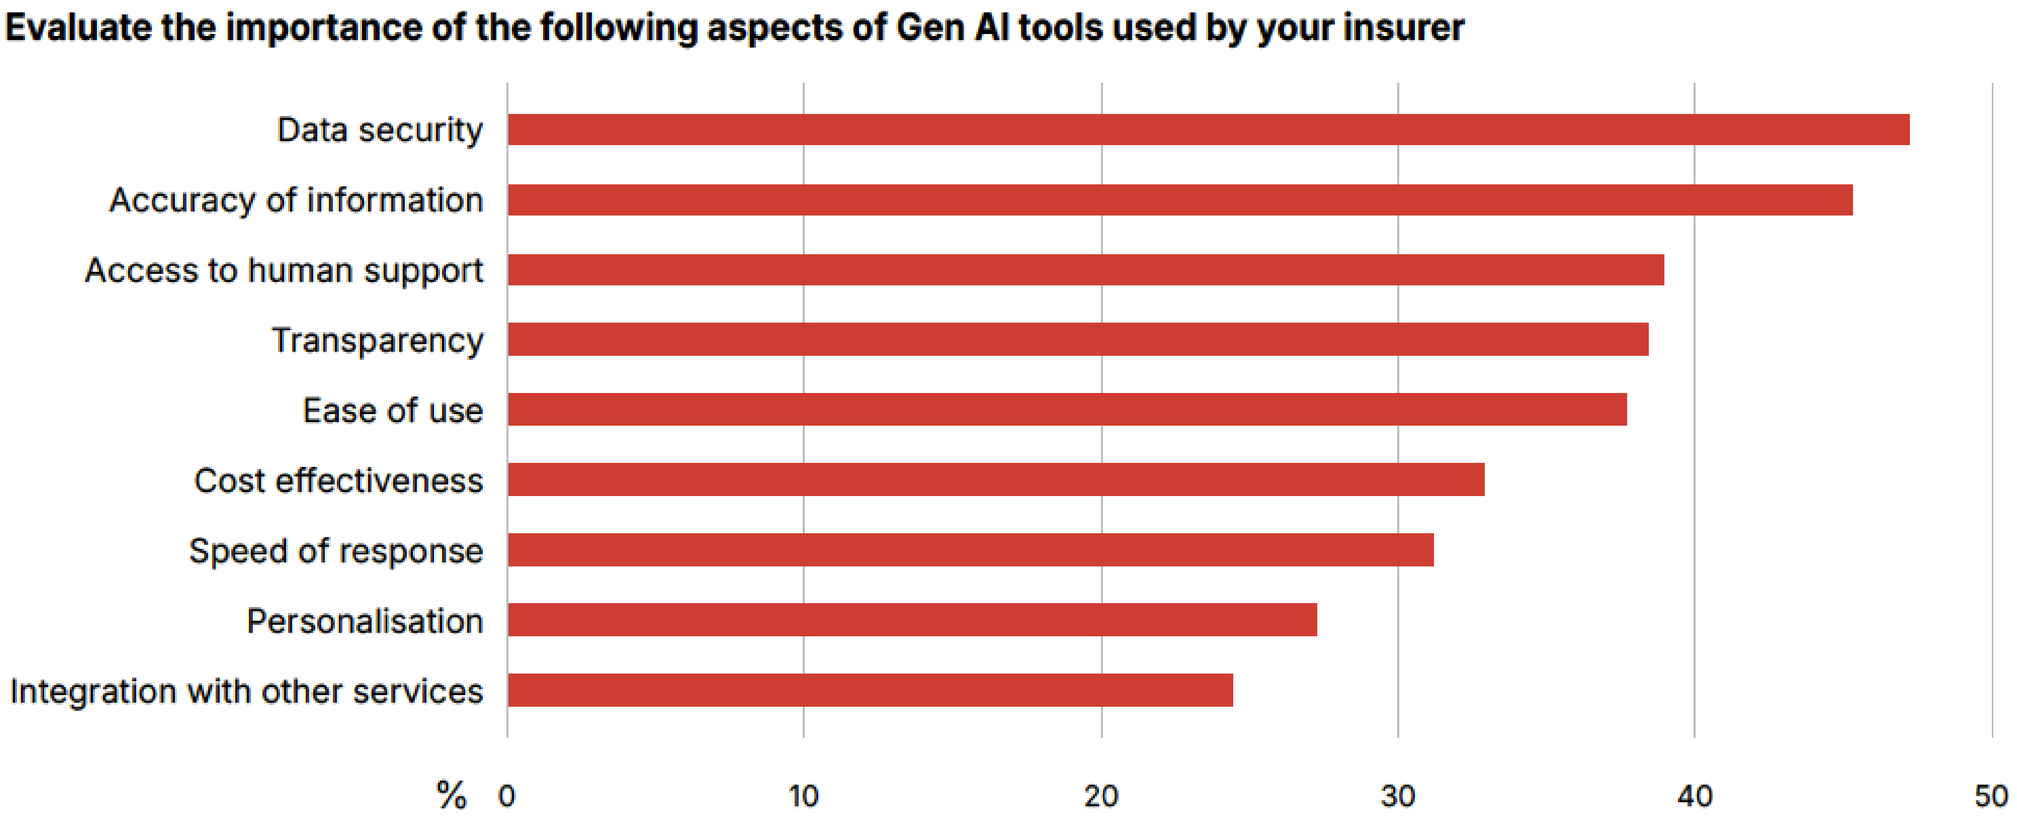

The Independent Use of Gen AI

68% of respondents have used tools like ChatGPT or Copilot when purchasing insurance.

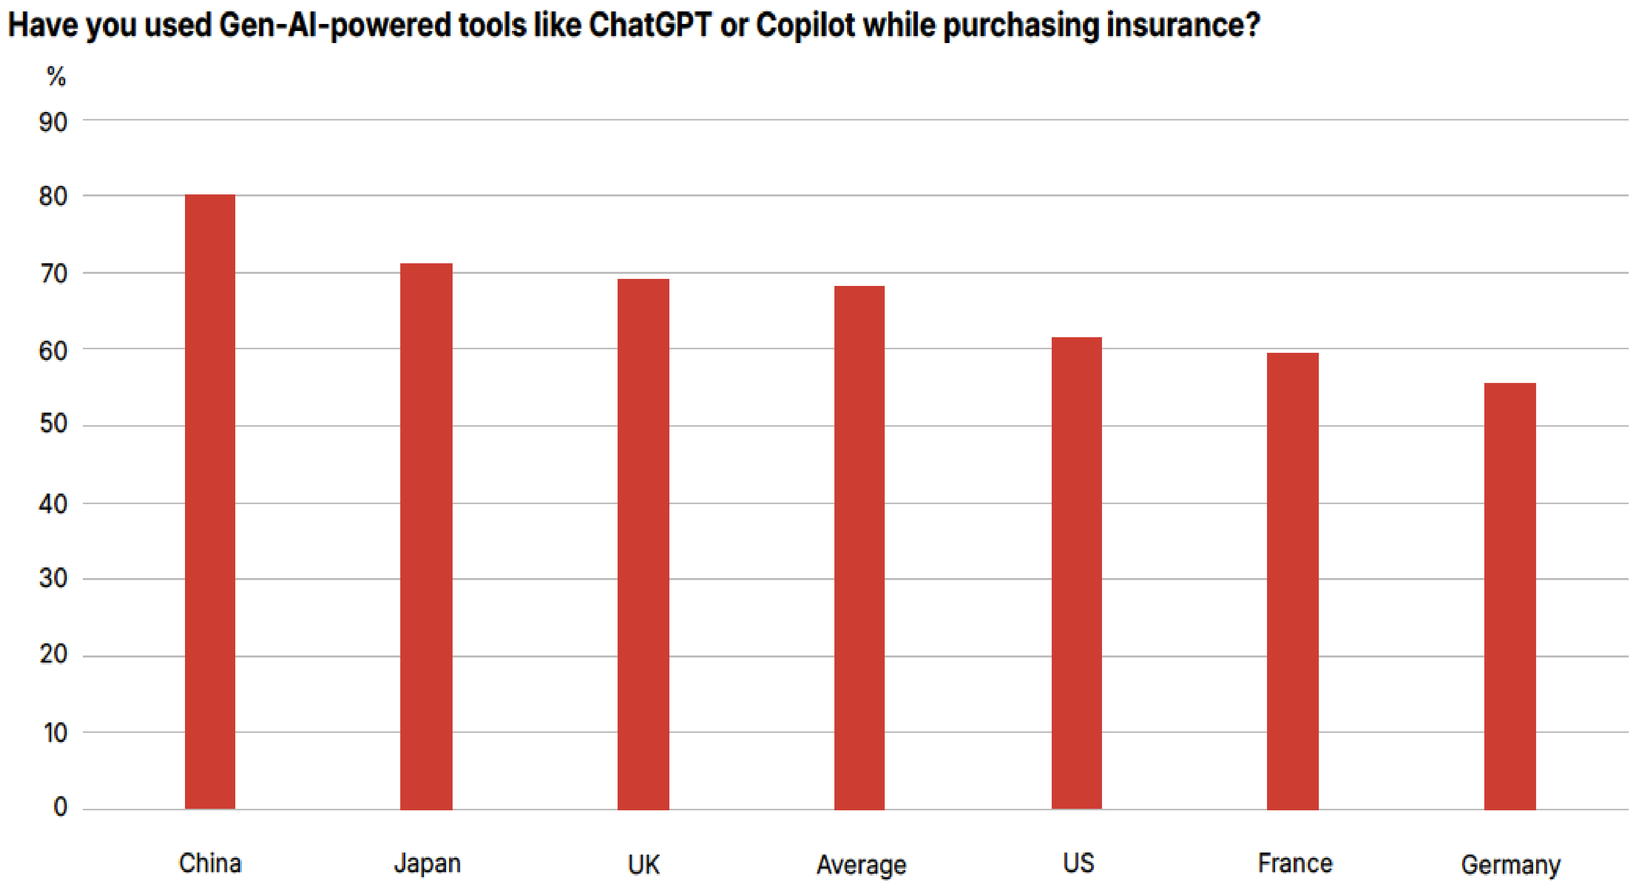

Source: Geneva Association insurance customer survey

Empowered Customers Change Market Dynamics

Off-the-shelf AI reduces:

- Information asymmetry
- Search costs
- Dependence on agents and brokers

Customers become:

- Better prepared
- More analytical
- More demanding

Reducing Information Asymmetry

Customers compare products and decode complex language in real-time.

“Confusopoly” is dead. Products must be transparent to withstand AI scrutiny.

Concerns and Challenges

	Top 3 concerns	
Insurer-provided Gen AI	Lack of human touch	▪ top concern (nearly 40% respondents), more pronounced in the Europe and the U.S
	Privacy and security	▪ 35% respondents unwilling to share personal data
	Accuracy and transparency	▪ The “black box” nature and “hallucinations”
Independent use of off-the-shelf Gen AI	Quality and reliability	▪ Model bias, reliability of information and advice generated
	Privacy	▪ Misuse or leakage of sensitive personal data leaking.
	Cybersecurity risks	▪ Phishing, identity theft, and hidden vulnerabilities

Human Touch

Data Privacy

Accuracy

Transparacy

The Hallucination Risk and Mitigation Strategies

Gen AI is probabilistic, not deterministic.

Risks:

- Incorrect policy interpretation
- Fabricated exclusions
- False confirmation of coverage

Consequences:

- Legal liability
- Regulatory penalties
- Reputational damage

RAG (Retrieval-Augmented Generation)

User Query → Retrieve Verified Doc (Policy) → Generate Answer based ONLY on Doc → Accurate Output

Geneva Association Gen AI and insurance report series

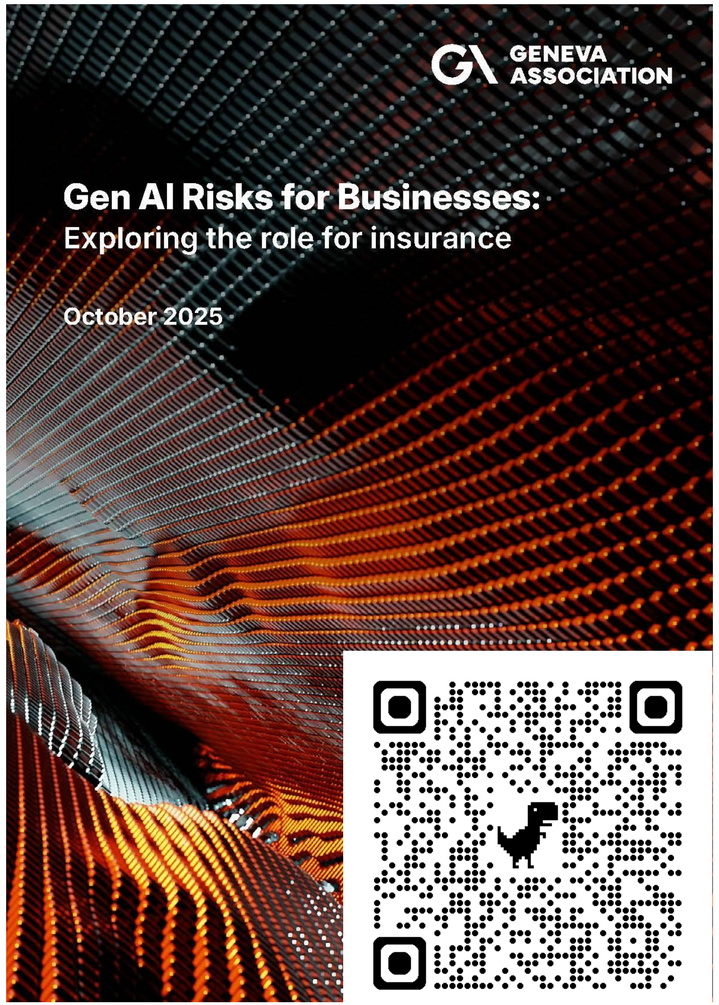

Report 1: [Gen AI Risks for Businesses: Exploring the role for insurance] (published on 2 October, 2025)

Report 2: [Gen AI in the Insurance Customer Journey]

(published on 20 November, 2025)

Rapid Adoption by Businesses

Adoption is accelerating globally.

- 71% of firms use Gen AI in at least one function (2025)
- Up from 33% in 2023

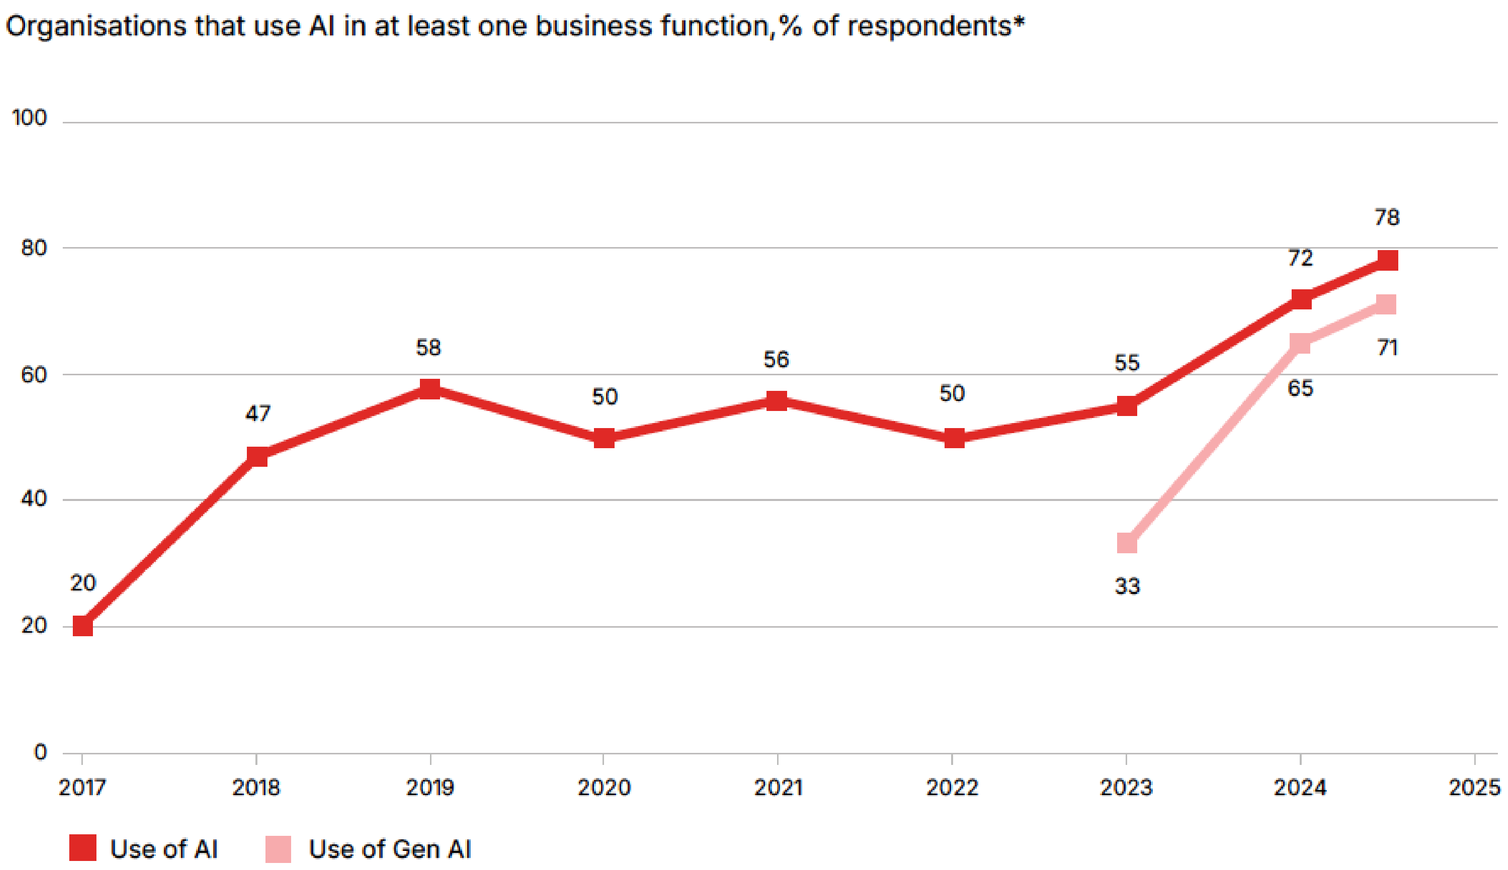

* In 2017, the definition for AI use was using it in a core part of the organisation's business or at scale. In 2018–19, it was embedding at least one AI capacity in business processes or products. Since 2020, it is that the organisation has adopted AI in at least 1 function.

Source: McKinsey[6]

Seven risk categories identified for Gen AI

- Operational, privacy & cybersecurity, ethical, regulatory, reputational, workforce, ESG

	Risk Category	Specific Risk
First-party operational risk	Operational	Algorithmic Errors; Stability; Reliability Black-Box Issues Malicious Attack
	Cybersecurity & Privacy	AI-driven Cyberattacks Data Privacy Violations
	Reputational & Market	Customer Trust and Brand Image Dependency & Competitive Risk
	Workforce Challenges	Job Displacement AI Skill Requirements
First-party operational risk +Third-party product risk	Regulatory & Compliance	Evolving AI Regulations Accountability and Liability
	Bias & Ethical Concerns	Discrimination and Bias Ethical Decision-Making
	ESG	Environmental and Energy

How Firms Use Gen AI

For what purposes do you use Gen AI at the company/organisation you work for or own?

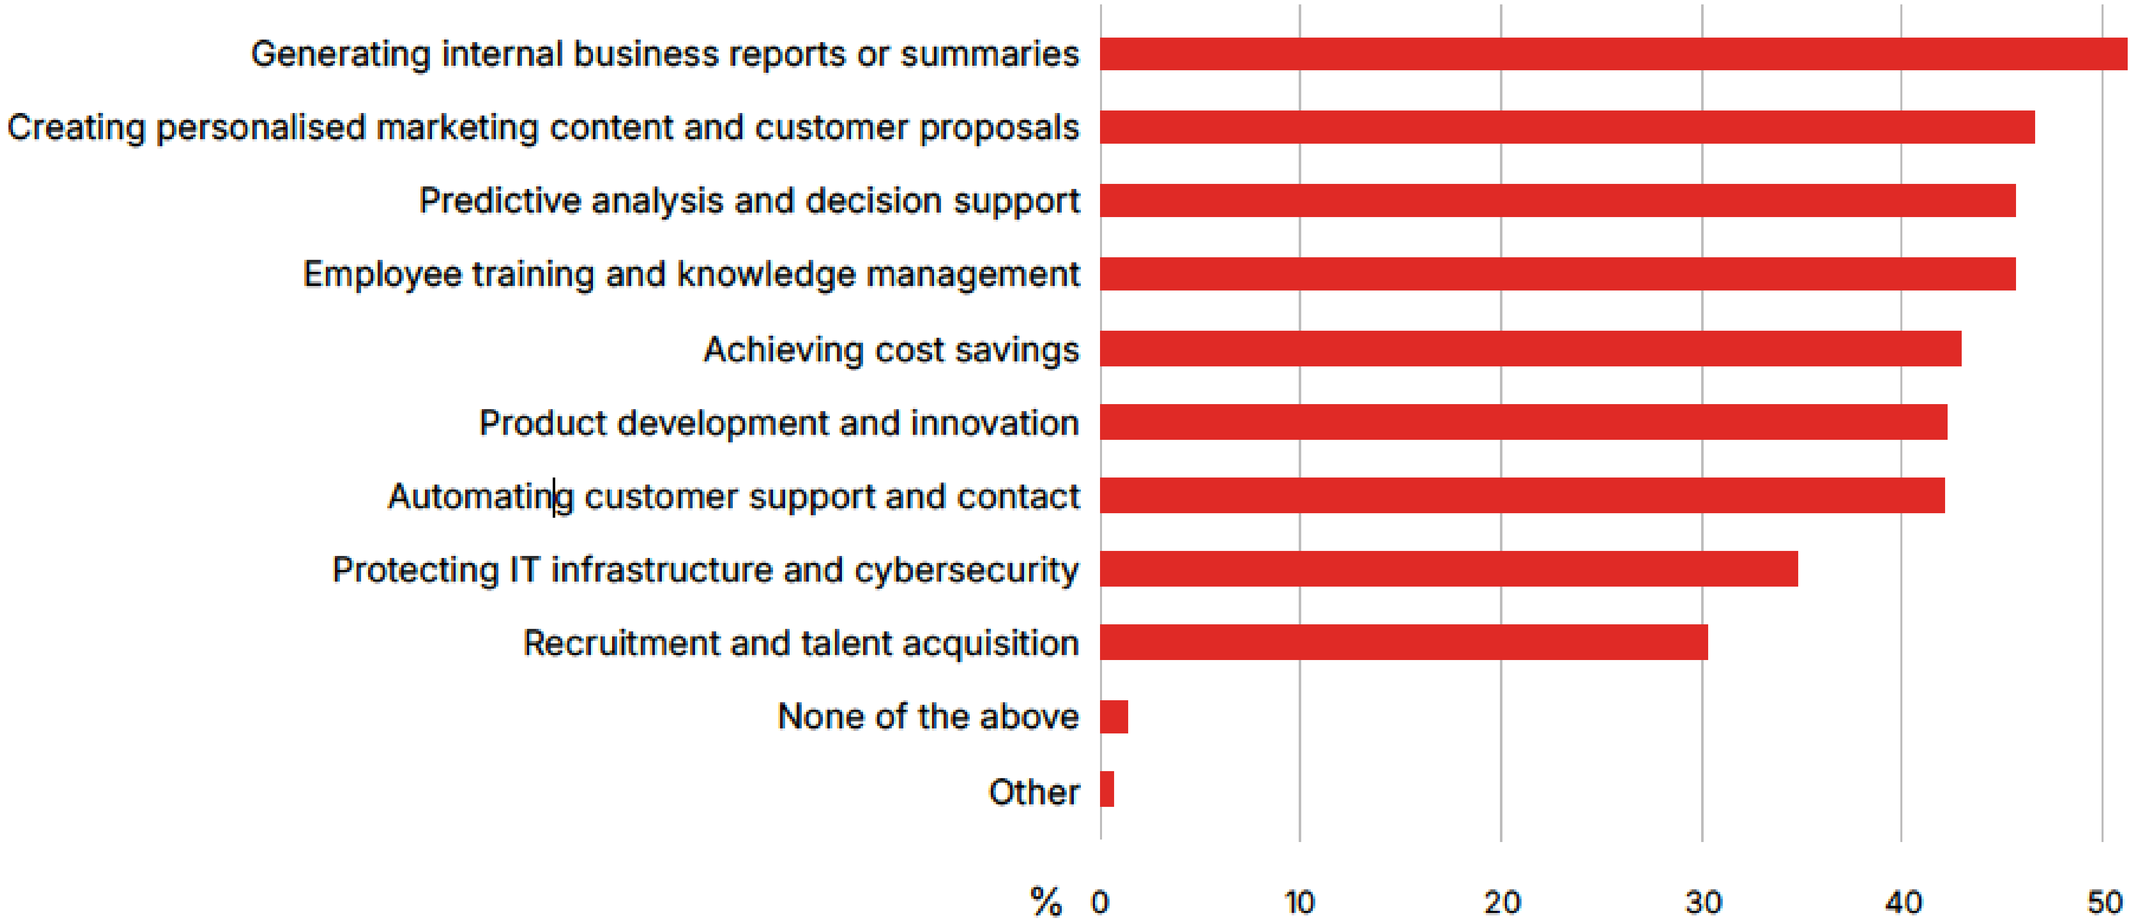

Source: Geneva Association business insurance customer survey

Implementation Challenges

What challenges have you encountered/do you anticipate when implementing Gen AI at the company/organisation you work for or own?

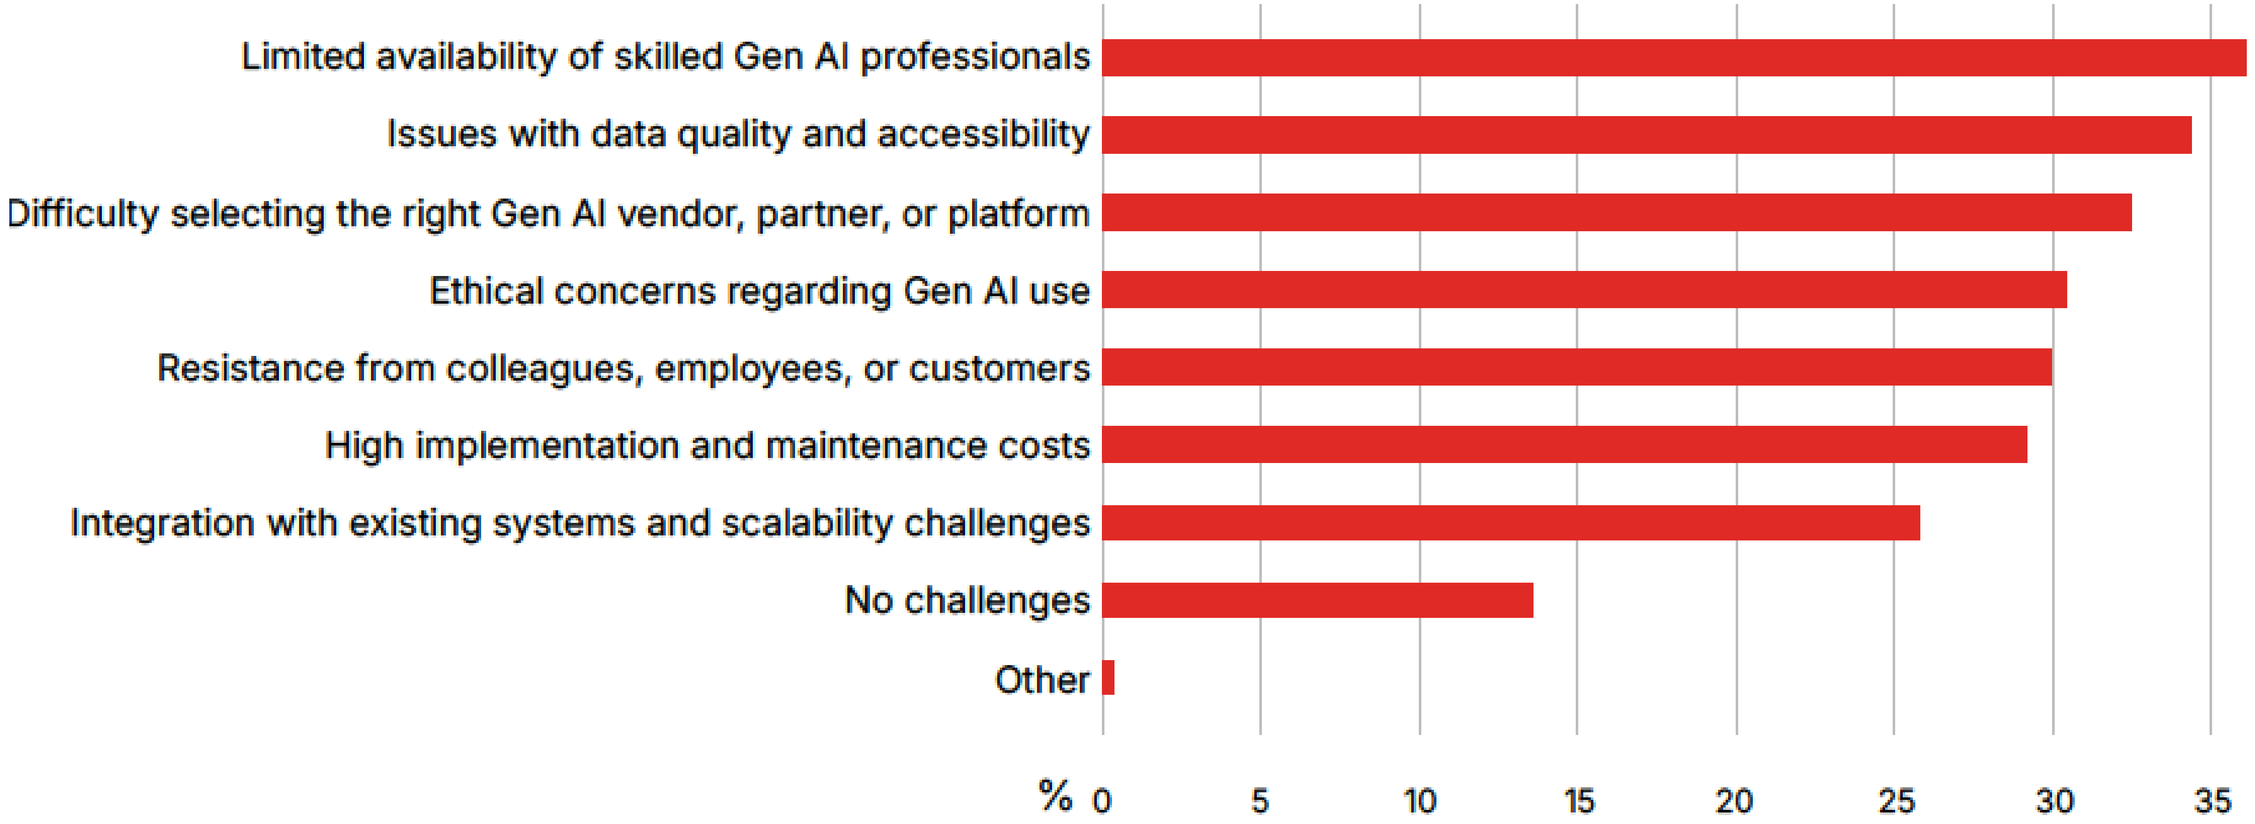

Source: Geneva Association business insurance customer survey

Failures & Incidents

What type(s) of failures or issues have you experienced or heard about when using Gen AI in the company/organisation you work for or own?

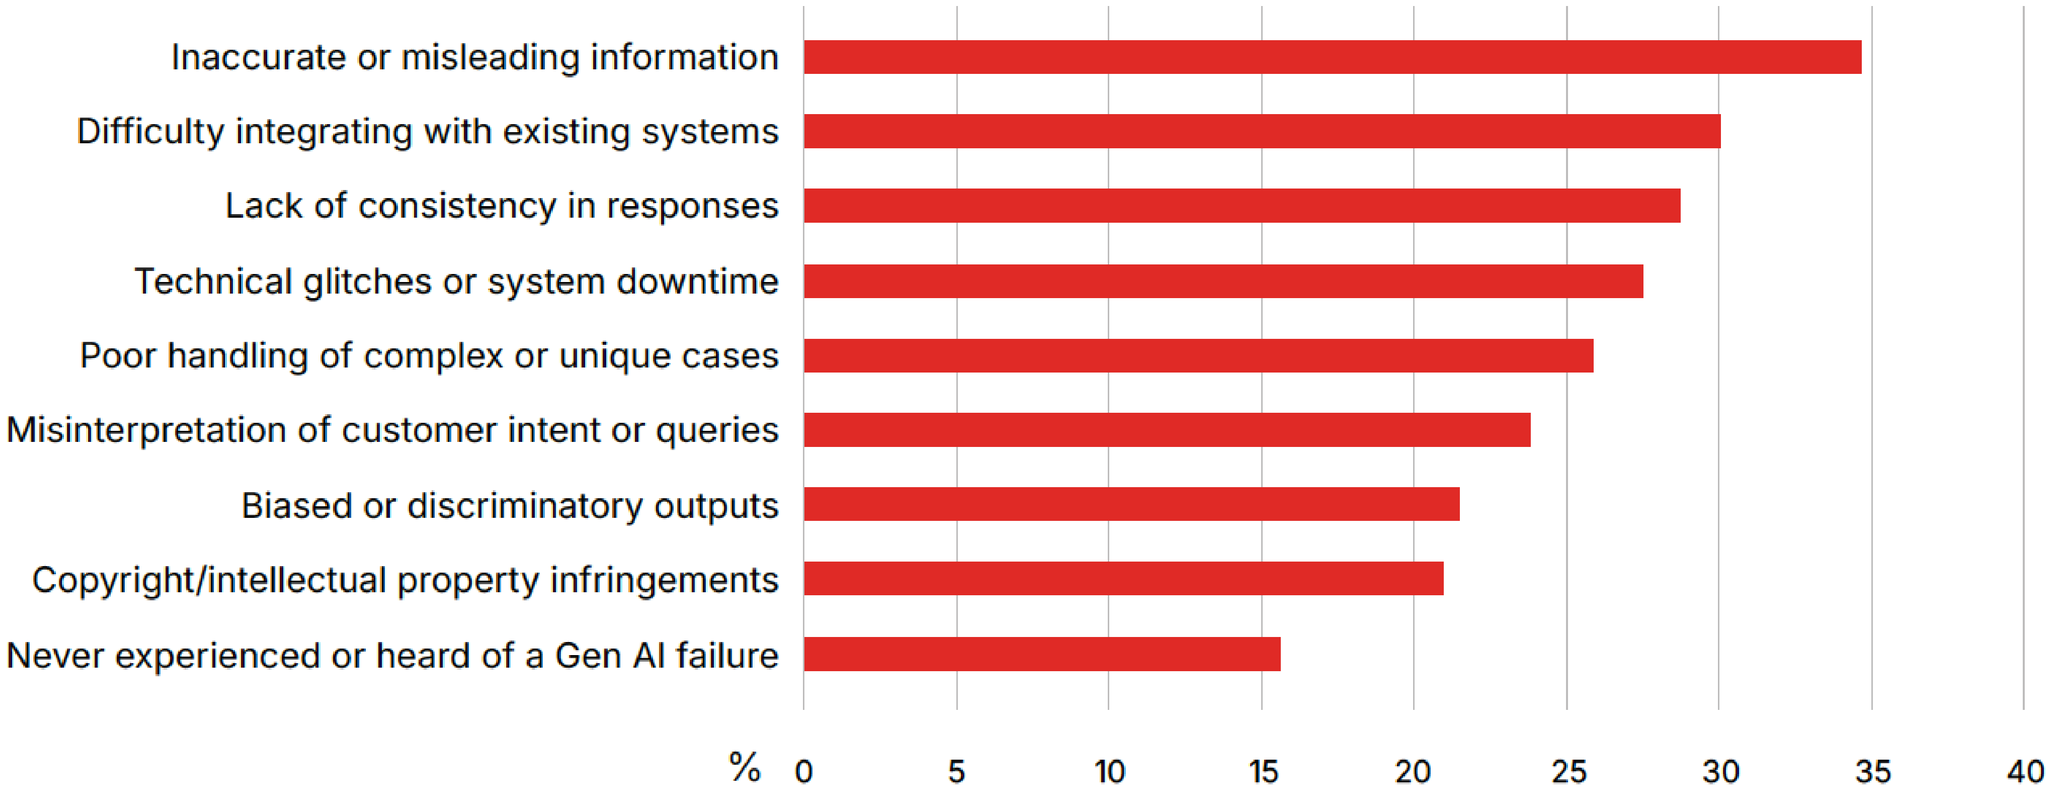

Source: Geneva Association business insurance customer survey

Demand for Gen AI Insurance

FIGURE 7: GEN AI RISKS BUSINESSES WANT TO INSURE

What Gen-AI-related risks do you want covered by your insurance?

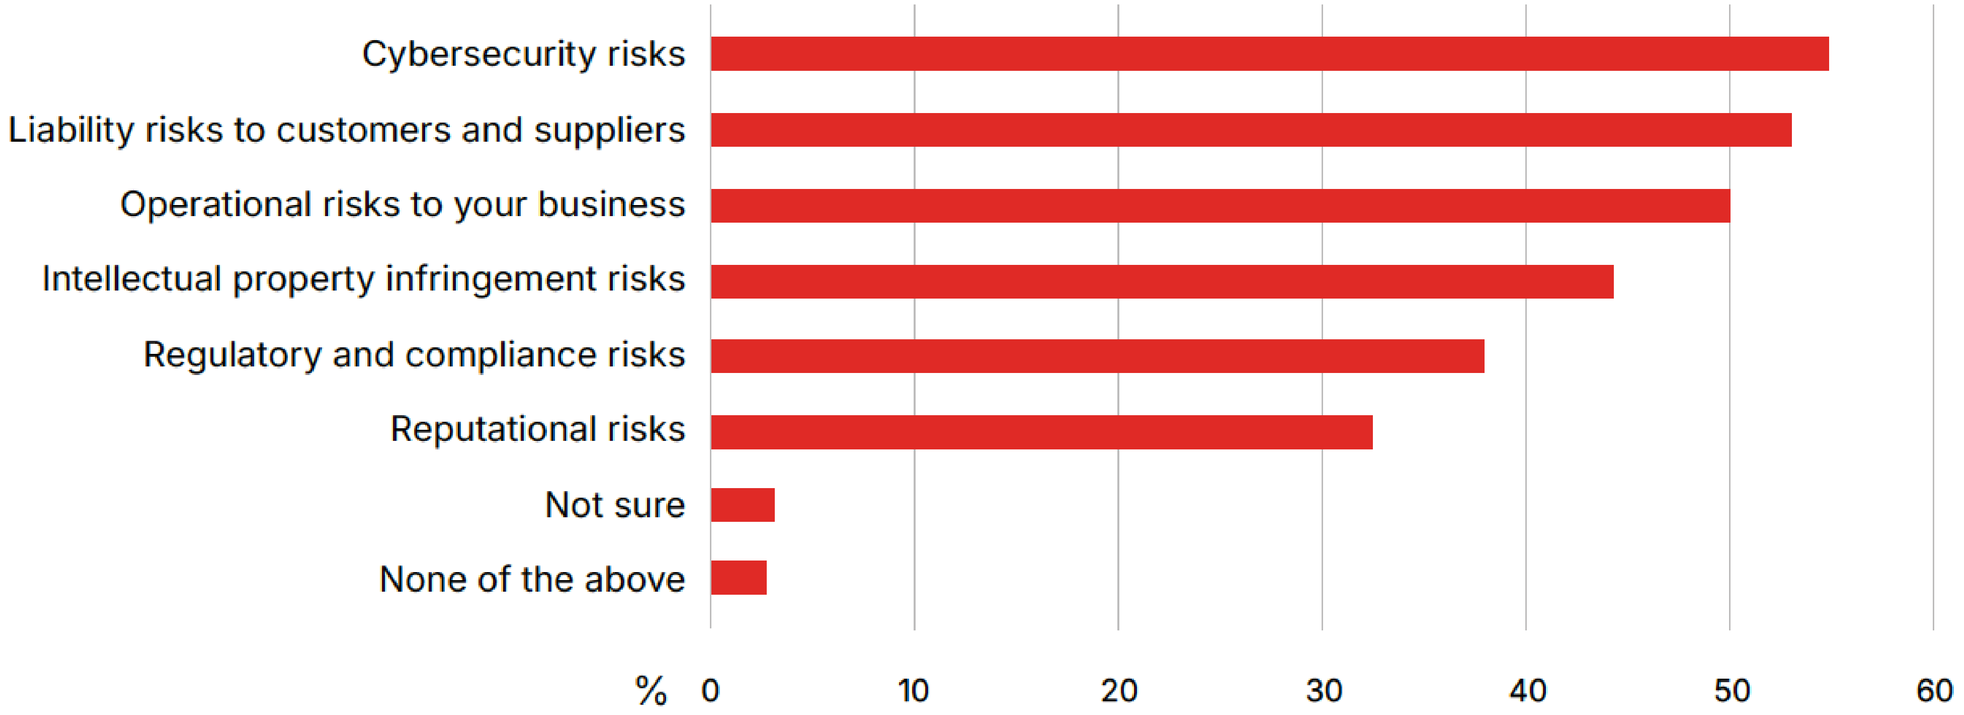

Source: Geneva Association business insurance customer survey

Demand for Gen AI Insurance

- Demand is strongest in the technology and finance sectors, among medium-to-large firms, and in markets with high Gen AI adoption
- Firms with high Gen AI exposure and severe past failures show higher insurance demand, suggesting potential adverse selection

From whom is Gen AI insurance in high demand

Business size

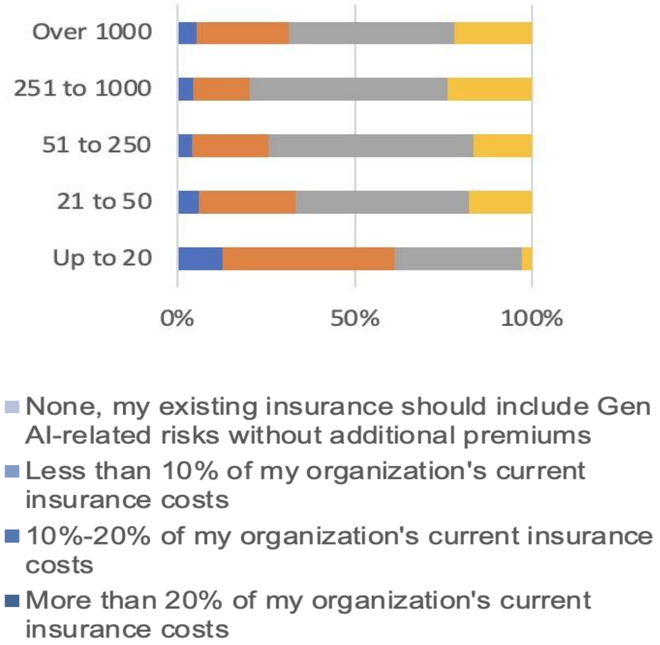

Industries

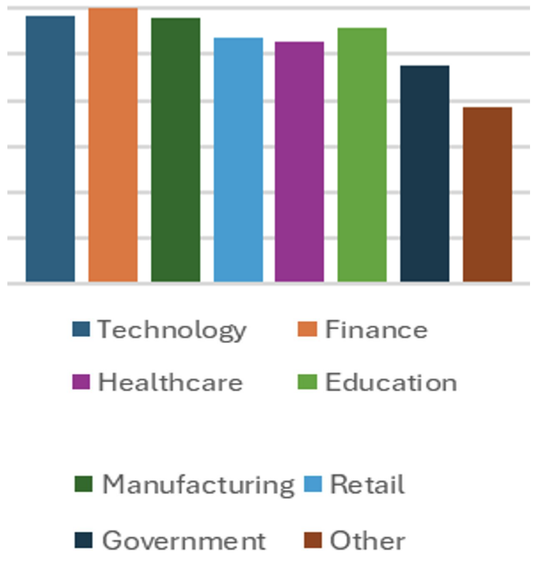

Across markets

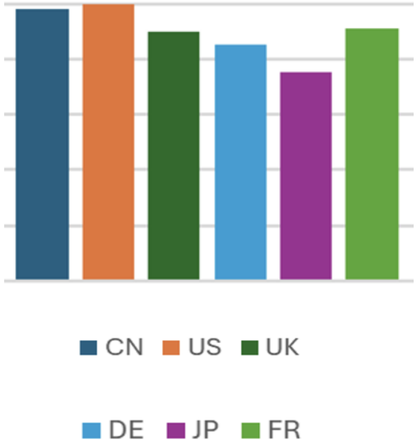

Risk exposure

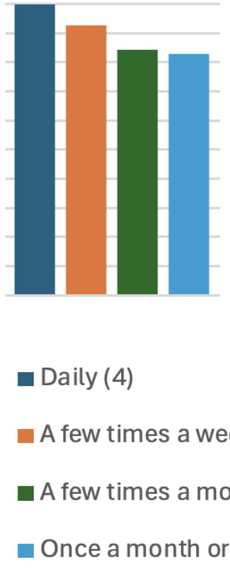

Failure severity

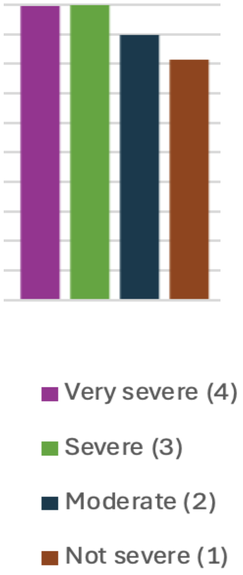

Willingness to Pay

FIGURE 8: WILLINGNESS TO PAY FOR GEN AI INSURANCE

How much extra would you be willing to pay for Gen-AI-related coverage?

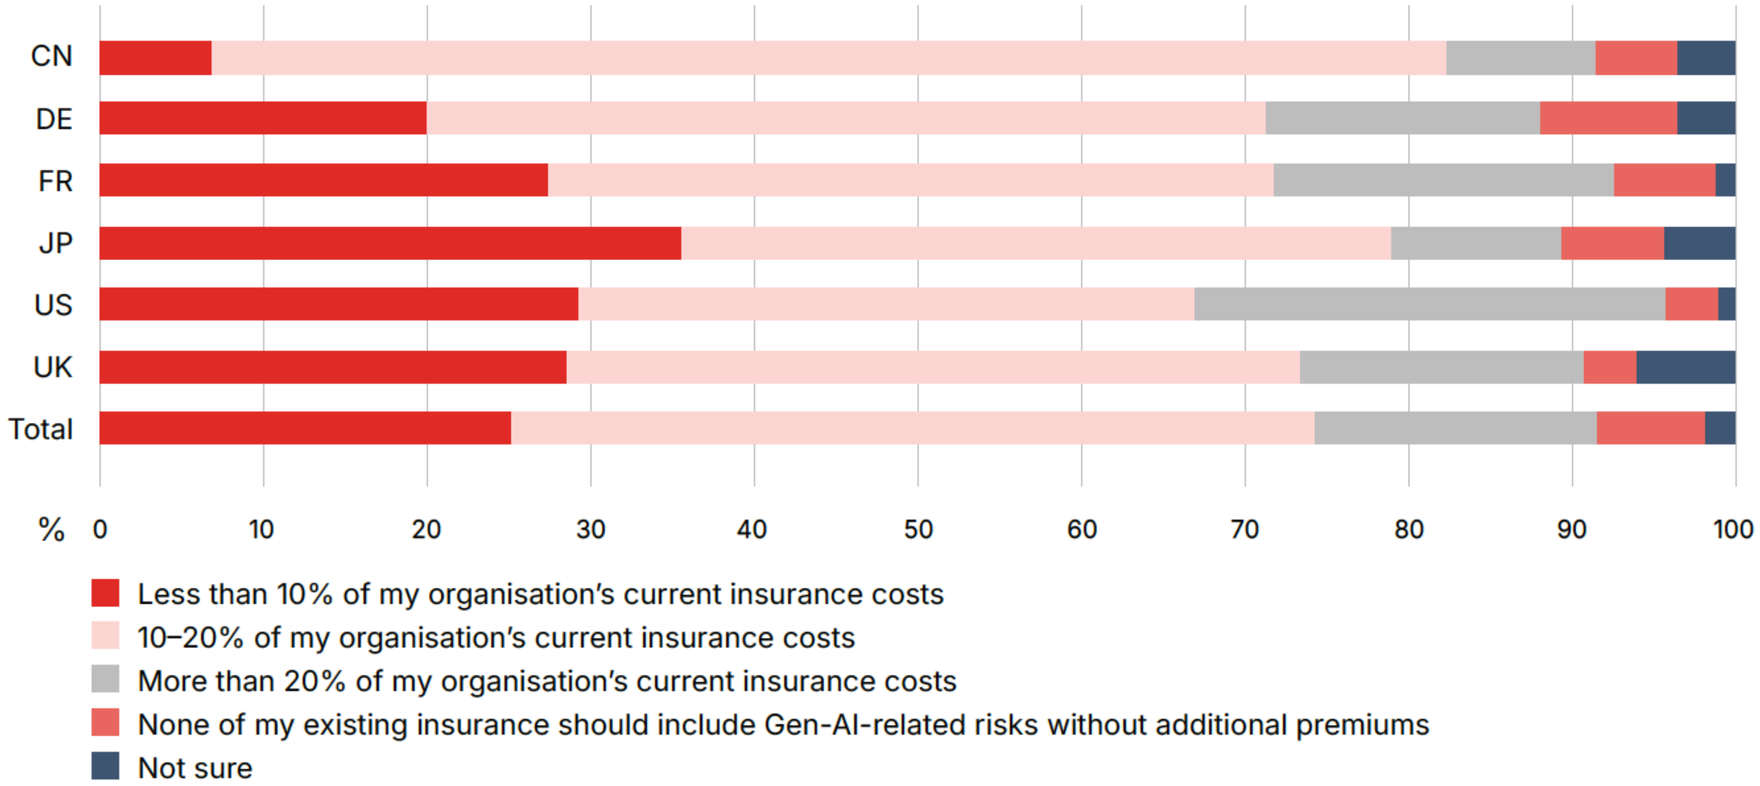

Source: Geneva Association business insurance customer survey

Supply of Gen-AI-related insurance

Insurability of Gen-AI-related risks

- Critical challenges for Gen AI:
  - Maximum possible loss
  - Average loss amount
  - Information asymmetry

- Mirrors the early evolution of cyber insurance

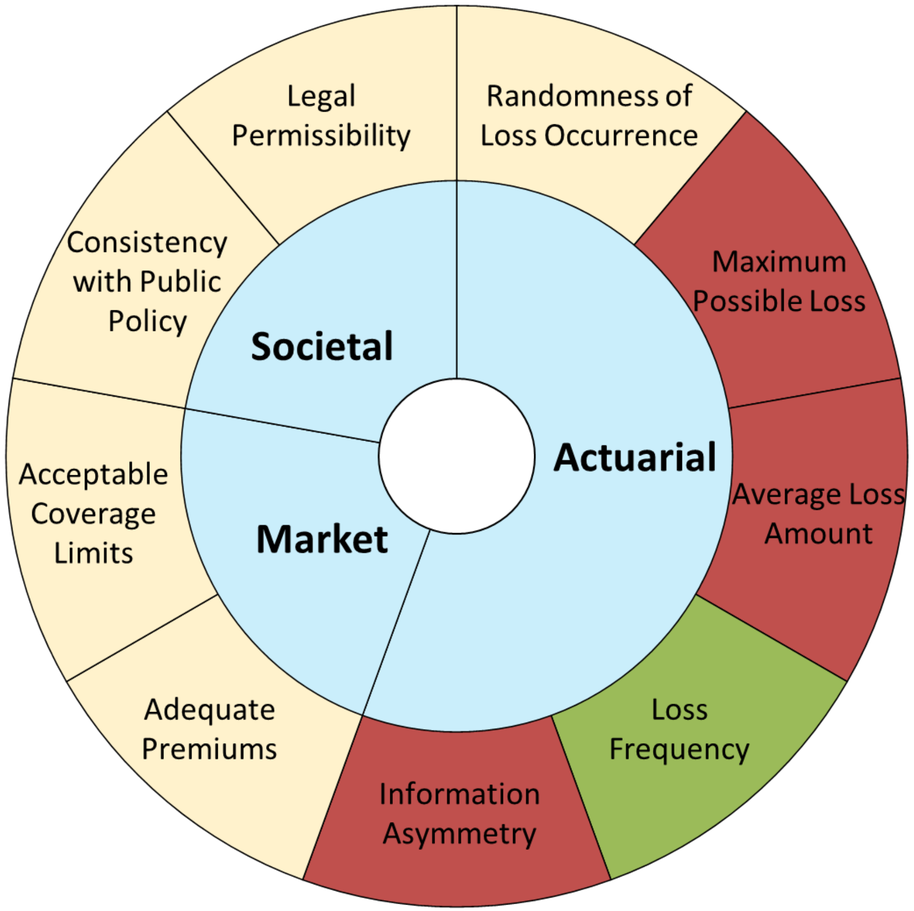

Source: Geneva Association

Emerging insurance solutions

- Gen-AI-related risks and associated coverage are still at an early stage but are expected to expand
- Effective Gen AI risk management requires a broad toolkit – a balanced approach combining risk avoidance, mitigation, retention, and transfer

Type of coverage	Description
Cyber insurance	Covers losses due to enhanced vulnerability of IT systems, as businesses and attackers use Gen AI tools, including data breaches, hacking, or model manipulation.
Professional liability (E&O)	Protects against claims arising from errors in Gen AI outputs, like the generation of misleading or harmful information.
Directors' & officers' liability	Protects directors and officers from legal action due to Gen-AI-related decisions or oversight failures.
Intellectual property	Protects against claims related to Gen AI's use of copyrighted or patented materials without permission.
Product liability	Covers claims due to harm caused by Gen AI outputs (e.g. misinformation, discriminatory content) or failure to perform as expected.
Dedicated, standalone AI insurance	Comprehensive, standalone coverage bundling multiple Gen-AI-specific exposures into a single policy.

II. Risk Analysis in the Technology Sector

PICC 中国人民保险

Three Stages of Technology Innovation

A R&D Stage ▶ B Technology Transfer Stage ▶ C Industrialization and Scale-up Stage

R&D risks refer to risks arising when technology companies, research institutions, universities and other technology actors form new technological achievements through basic research, applied research, or experimental research.

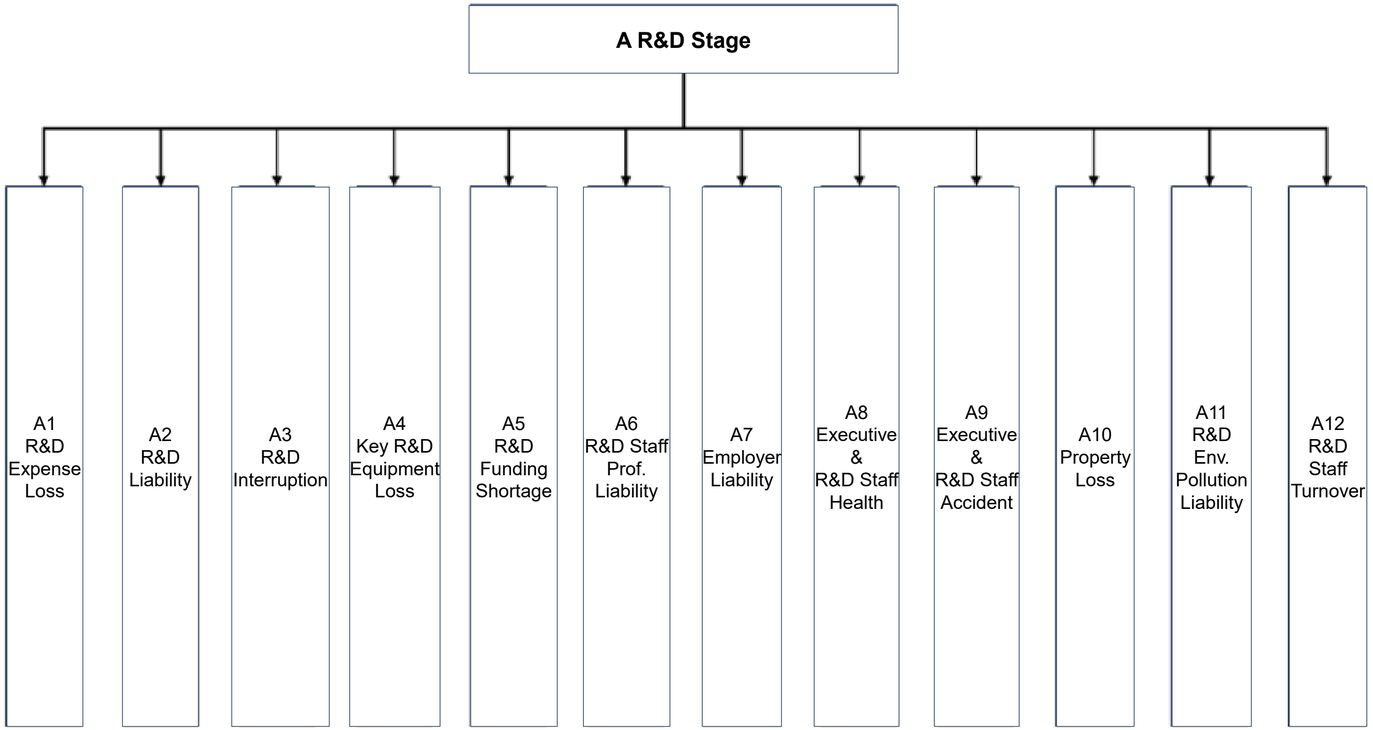

Risks Related to the R&D Stage

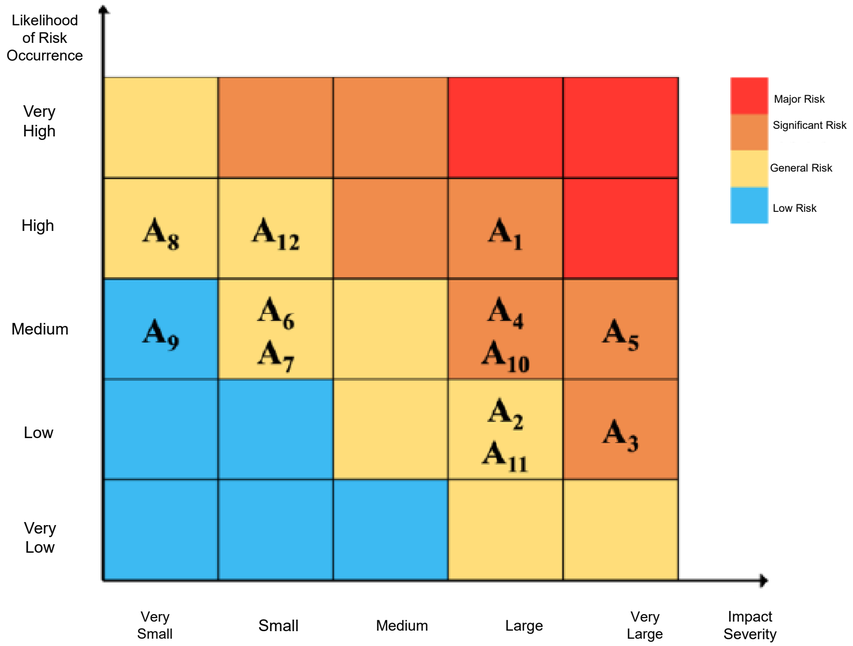

Risk Matrix for the R&D Stage

II. Risk Analysis in the Technology Sector

PICC 中国人民保险

Three Stages of Technology Innovation

A R&D Stage ▶ **B Technology Transfer Stage** ▶ C Industrialization and Scale-up Stage

Technology transfer risks refer to risks arising when technology companies, research institutions, universities and other technology actors transform technological achievements into products or services with practical application value and commercialization potential.

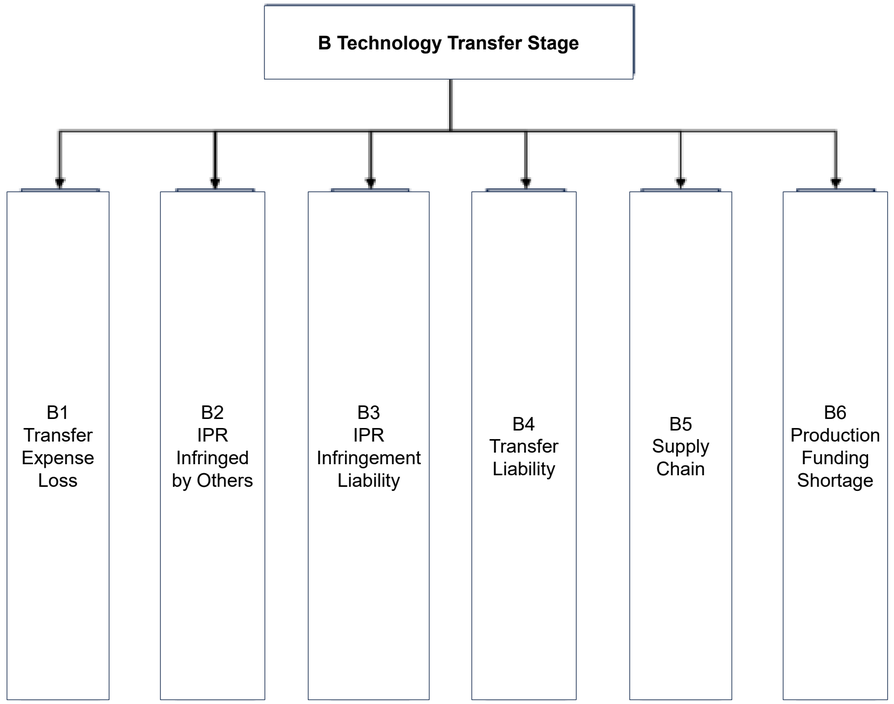

Risks Related to the Technology Transfer Stage

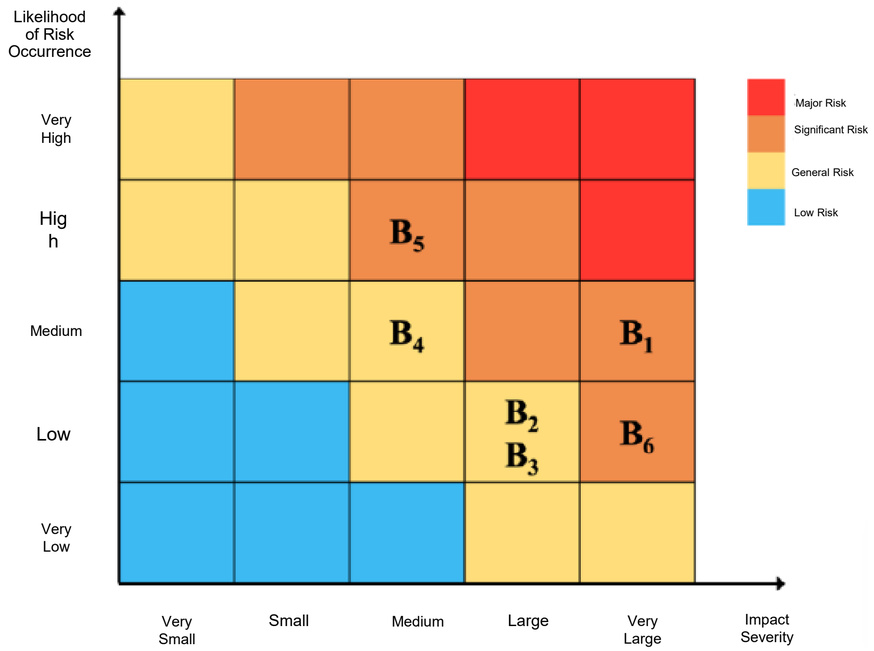

Risk Matrix for the Technology Transfer Stage

II. Risk Analysis in the Technology Sector

PICC 中国人民保险

Three Stages of Technology Innovation

A R&D Stage ▶ B Technology Transfer Stage ▶ **C Industrialization and Scale-up Stage**

Industrialization and scale-up is the key link through which technology innovation realizes its value.

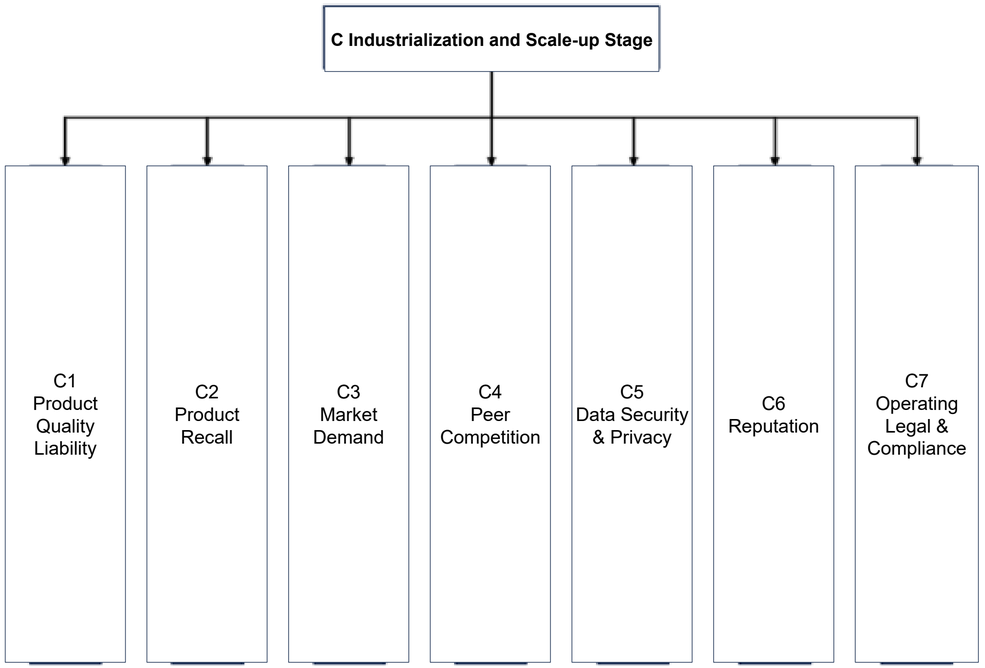

Risks Related to the Industrialization and Scale-up Stage

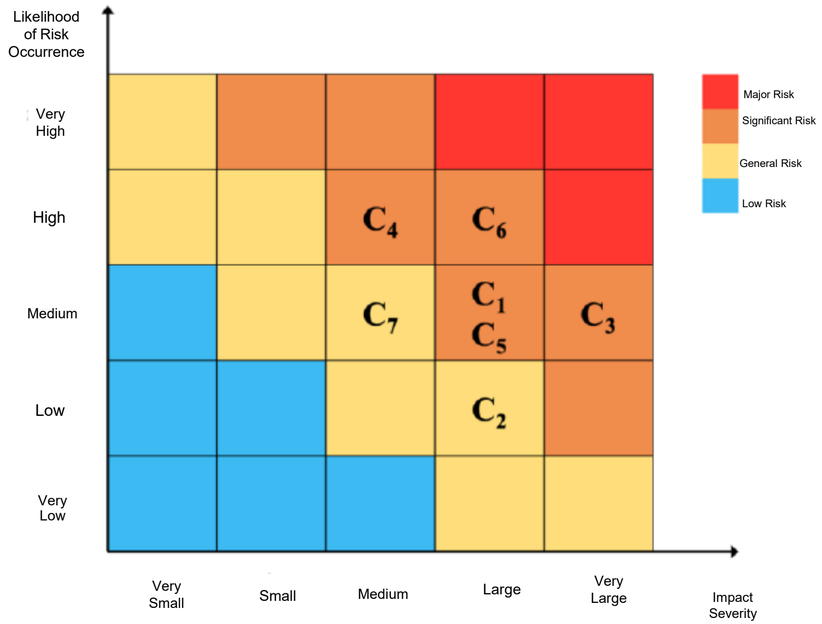

Risk Matrix for the Industrialization and Scale-up Stage

Insurance for Infringement Liability Arising from Generative AI-Generated Content

Enterprises, public institutions, and other organizations that use generative AI technology to provide generative AI services may act as the policyholder and insured under this insurance contract.

Insurance Liability

- Where, in the course of providing generative AI services, the insured, due to fault or negligence, causes the generated content output based on user input to infringe a third party's lawful rights and interests, including copyright, portrait rights, and reputation rights, and the relevant rights holder first requests handling by the competent authority, files a lawsuit accepted by a court, or files an arbitration request accepted by an arbitration institution during the insurance period, the insurer shall indemnify the insured for the following costs and economic compensation liability in accordance with this insurance contract:
- Arbitration or litigation costs and other necessary and reasonable legal expenses paid with the insurer's prior written consent.
- Economic compensation liability that the insured is legally required to bear.

Ruo (Alex) Jia 贾若
ruo.jia@pku.edu.cn
alex_jia@genevaassociation.org

Driving InsurTech Innovation with Real-World Insurance Data: From ML to AI

Real-world Usecases

Zhiyu (Frank) Quan

University of Illinois Urbana-Champaign

Overview

- InsurTech Innovation
- Data is the New Oil: Federated Learning
- Automated Machine Learning
- Agentic AI: Claim Automation
- AI is neither a Hype nor Magic

InsurTech Innovations with Real-World Data

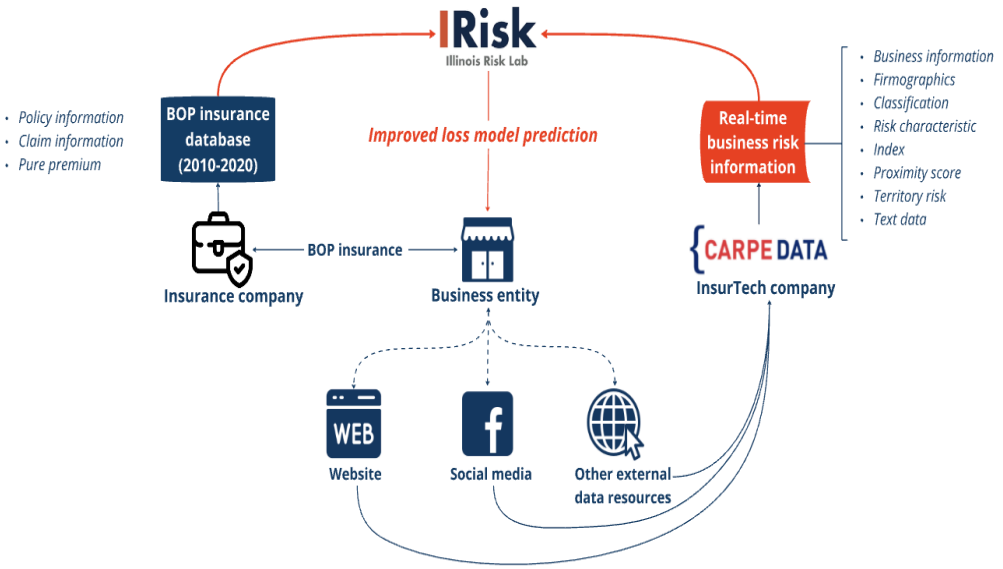

- The evolving risk landscape calls for ***additional*** predictive factors beyond ***traditional*** rating variables to achieve more accurate pricing[1].
- The ***university*** serves as an independent third-party evaluator for the InsurTech data vendor, quantifying the predictive lift achieved when insurers incorporate additional predictive factors into their in house models.

[1] Quan, Z., Hu, C., Dong, P., & Valdez, E. A. (2025). Improving business insurance loss models by leveraging InsurTech innovation. *North American Actuarial Journal*, 29(2), 247-274.

Federated Learning: Facilitates Industry-wide Partnership

- Can we solve industry-wide ***pain points*** while preserving ***data privacy***?[2]
- Even large carriers can gain valuable insights from data contributed by smaller insurers.

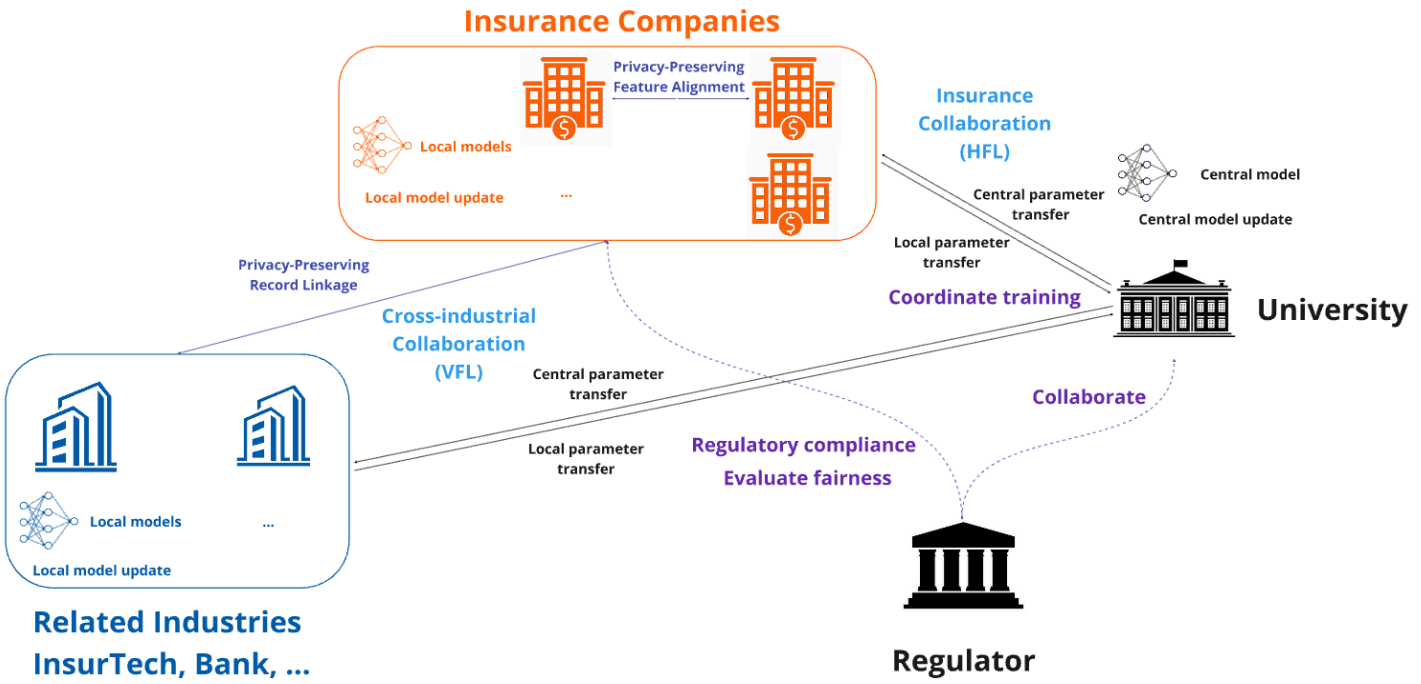

[2] Dong, P., Feng, R., Quan, Z., & Wang, T. (2026). Bridging the Divide While Walking a Tightrope: Evidence from the Insurance Industry on Federated Data Sharing.

Automated Machine Learning

- ***Empower*** actuaries to unlock the power of data science through AutoML[3].

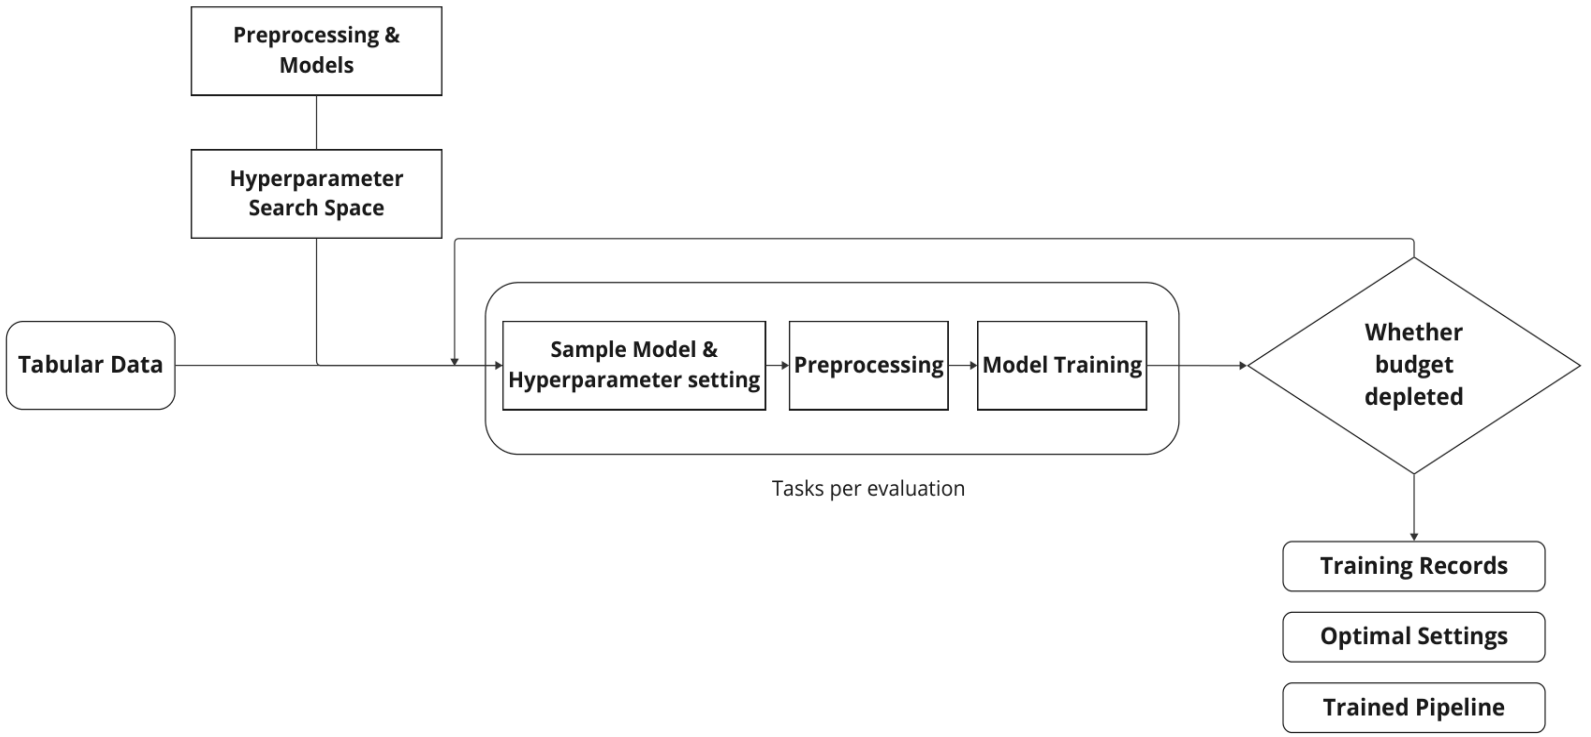

[3] Dong, P., & Quan, Z. (2025). Automated machine learning in insurance. *Insurance: Mathematics and Economics*, 120, 17-41.

- Do we still need data scientists/machine learning researchers in the AI era? ***Domain knowledge*** and ability to transform business problem into optimization frameworks are more important than ever[4].

- Interested to try out[5]?

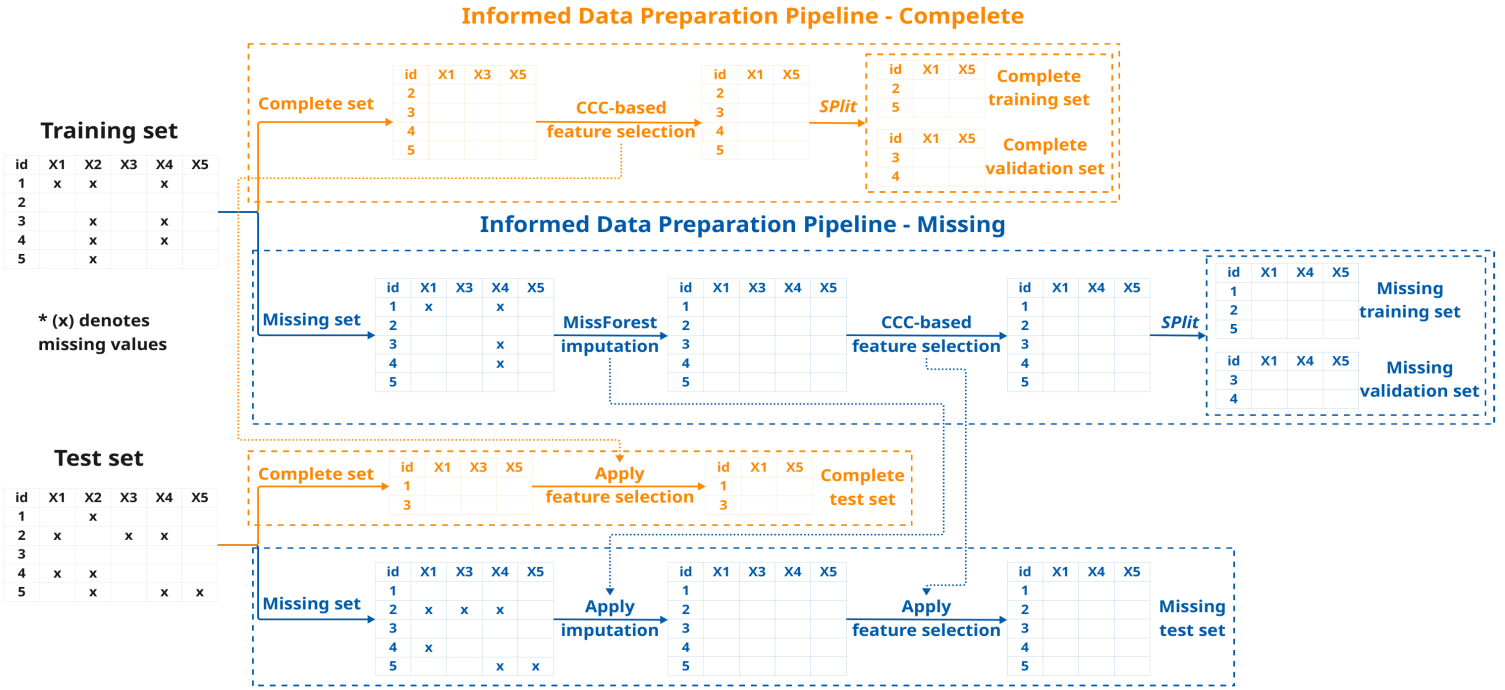

[4] Guo, J., Dong, P., & Quan, Z. (2026). Starting Off on the Wrong Foot: Pitfalls in Data Preparation.

[5] https://github.com/PanyiDong/InsurAutoML

Claim Agentic AI

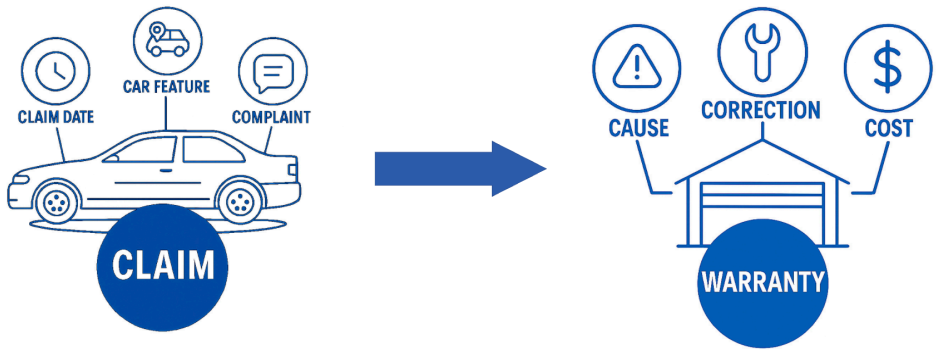

- Customer **complaint** descriptions — free-text narratives of vehicle issues.
- Diagnosed **causes** — brief technician notes or internal codes.
- **Corrective actions** — usually free-text descriptions of repairs.
- **Loss codes** — standardized categorizations of claim types.

- Large Language Models (LLMs) can ***assist*** claim adjusters, service writers, and customers by quickly and accurately predicting repair needs—parts, labor, and cost[6].

[6] Mo, Z., Quan, Z., O'Donohue, E., & Zhong, K. (2026). Claim Automation using Large Language Model.

Localized Open-Source LLM

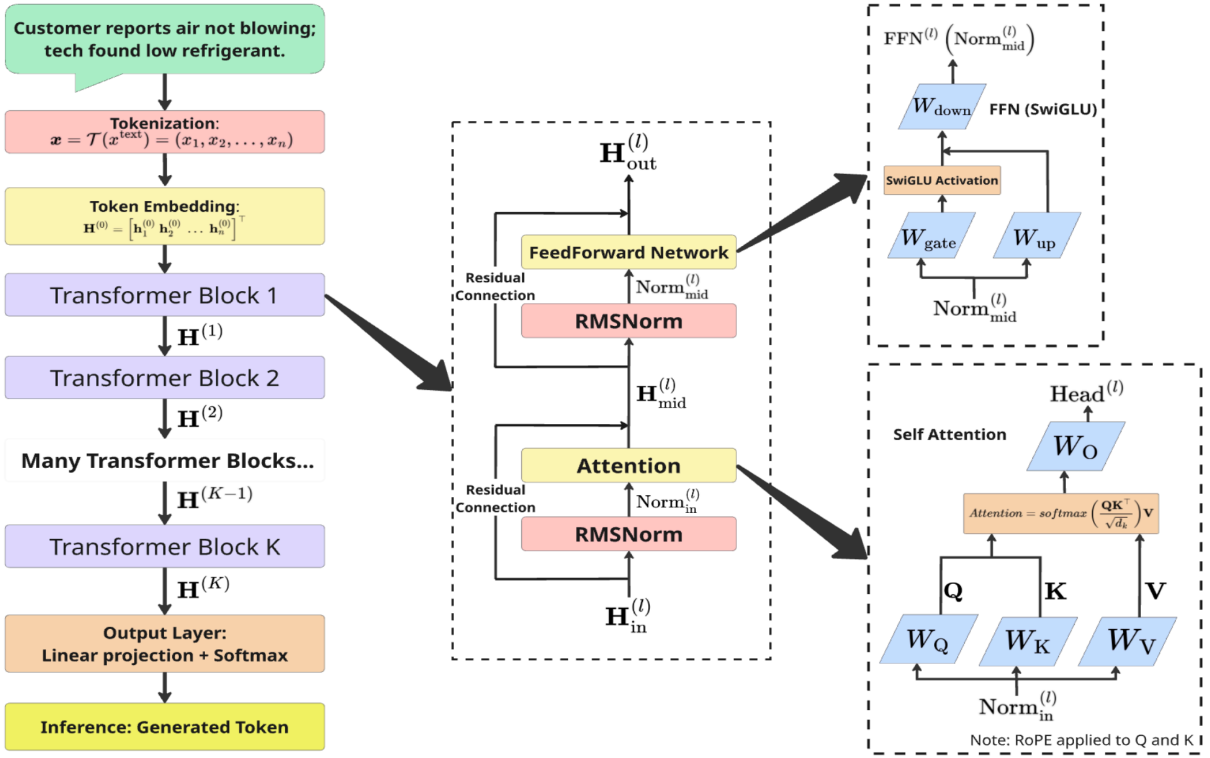

- LLMs can be ***locally*** deployed on GPUs on-prem. We fine-tune an ***open-source*** model (DeepSeek-R1-Distill-Llama-8B) on an 8×A100 GPU cluster with ***9 millions*** historical claim records.

Agent Handle Routine Claims

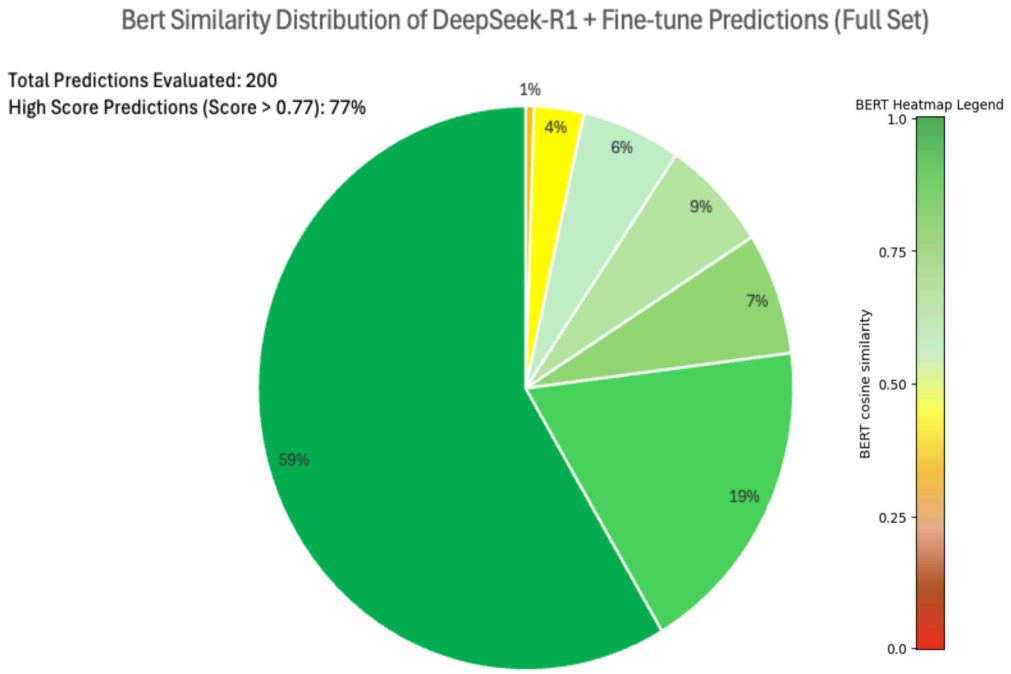

- Human and automated evaluations on model prediction:
  - **85.8% of predictions scored above 0.77**, meaning they are semantically correct with only minor wording differences.
  - Most remaining predictions fall between 0.4 and 0.8, reflecting partial correctness.
  - **A small portion (~2%) scored below 0.3**, indicating occasional misunderstandings or predicting a different corrective action than truth.

Commercial Sate-Of-The-Art LLMs

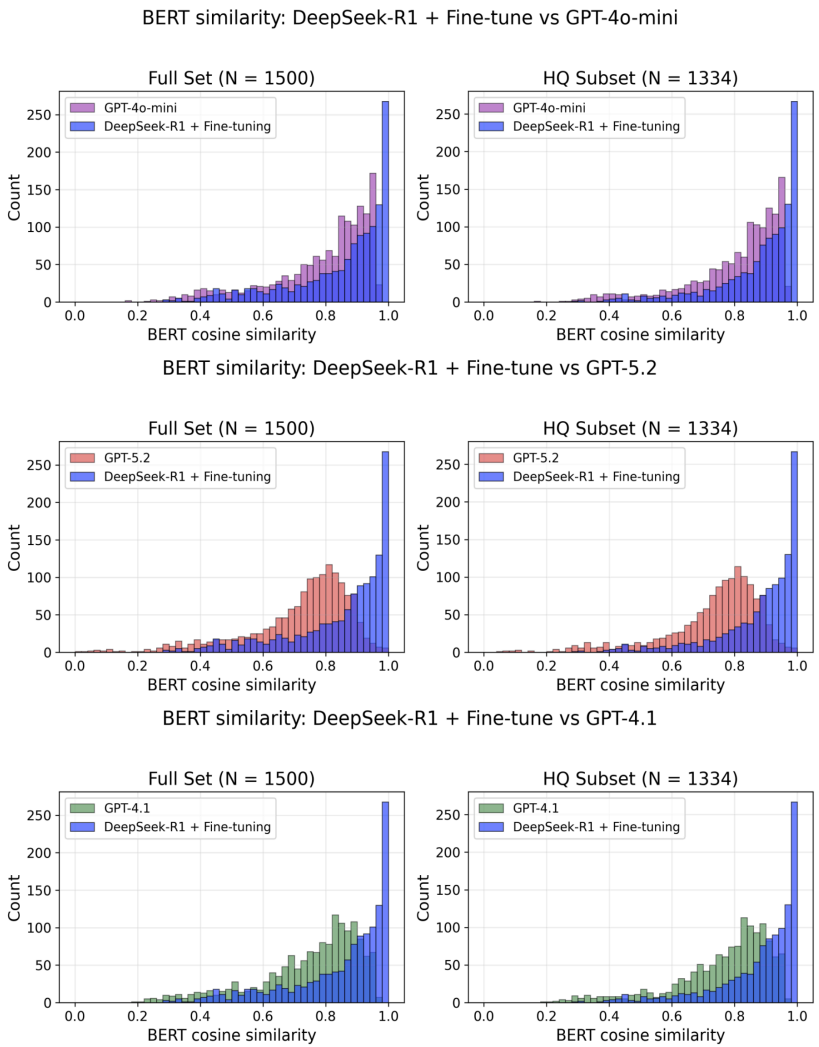

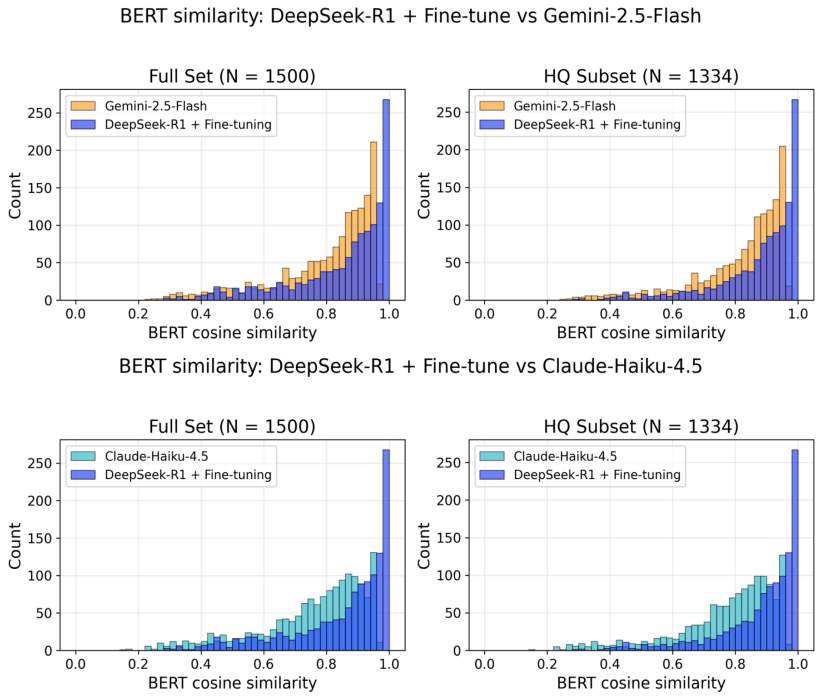

AI is neither a Hype nor Magic

- LLMs are not designed to be intelligent.
- Insurance data remains siloed.
- Insurers hold a wealth of proprietary data to unlock AI's potential.
- The competitive edge will belong to those who unlock its AI potential first.

Thank you! Q&A

Appendix A. Human Evaluation

Table 1: Human evaluation scoring guidelines used in all annotation tasks.

Score	Meaning	Description
1.0	Perfect Match	The predicted actions fully and precisely cover the ground-truth actions. No missing or incorrect actions.
0.8	Strong Match	The predicted actions mostly cover the ground truth, with only minor omissions or paraphrasing and no major errors.
0.6	Partial Match	The prediction covers some key actions but misses important parts. Still related and somewhat helpful.
0.4	Weak Match	Very limited overlap. Only a few actions (or fragments) are relevant to the ground truth.
0.2	Minimal Match	Prediction is largely irrelevant or incorrect, with almost no valid overlap with the ground truth.
0.0	No Match	Completely unrelated or invalid prediction, with no correspondence to the ground truth.

Appendix B. Automated Evaluation

C Formal Definitions of Automated Evaluation Metrics

In this section, we formally define all automated evaluation metrics used in section 3.4. Let y denote the reference corrective action and $\hat{y}$ the model-predicted action, both represented as natural language strings. All metrics quantify the similarity between $\hat{y}$ and y from different perspectives.

Edit distance similarity. Let $\mathrm{ED}(\hat{y}, y)$ denote the character-level Levenshtein edit distance between $\hat{y}$ and y. We define the normalized edit-distance similarity as

$$\mathrm{Sim}_{\mathrm{ED}}(\hat{y}, y) = 1 - \frac{\mathrm{ED}(\hat{y}, y)}{\max\left(|\hat{y}|, |y|\right)}, \tag{C.1}$$

where $|\cdot|$ denotes string length (number of characters).

N-gram precision. Let $\mathcal{G}_n(s)$ denote the multiset of n-grams (tokens) extracted from a string s, and let $c_s(g)$ denote the count of n-gram g in s. We define n-gram precision as

$$\mathrm{Prec}_n(\hat{y}, y) = \frac{\sum_{g \in \mathcal{G}_n(\hat{y})} \min\left(c_{\hat{y}}(g),\, c_y(g)\right)}{\sum_{g \in \mathcal{G}_n(\hat{y})} c_{\hat{y}}(g)}. \tag{C.2}$$

Sentence BLEU. We adopt sentence-level BLEU as implemented in sacreBLEU (Post, 2018), based on the original BLEU formulation (Papineni et al., 2002). The clipped n-gram precision is defined as

$$p_n(\hat{y}, y) = \frac{\sum_{g \in \mathcal{G}_n(\hat{y})} \min\left(c_{\hat{y}}(g),\, c_y(g)\right)}{\sum_{g \in \mathcal{G}_n(\hat{y})} c_{\hat{y}}(g)}. \tag{C.3}$$

Sentence BLEU is computed as the geometric mean of modified n-gram precisions with a brevity penalty:

$$\mathrm{BLEU}(\hat{y}, y) = \mathrm{BP} \cdot \exp\left(\sum_{n=1}^{N} w_n \log p_n(\hat{y}, y)\right), \tag{C.4}$$

where w_n are uniform weights, $N = 4$ in our experiments, and the brevity penalty is

$$\mathrm{BP} = \begin{cases} 1, & |\hat{y}| > |y|, \\ \exp\left(1 - \frac{|y|}{|\hat{y}|}\right), & |\hat{y}| \le |y|. \end{cases} \tag{C.5}$$

TF–IDF cosine similarity. Let $\mathcal{D} = \{y_1, \ldots, y_T, \hat{y}_1, \ldots, \hat{y}_T\}$ denote the combined corpus consisting of all reference and predicted actions. A TF–IDF vectorizer is fitted on $\mathcal{D}$ with n-gram range $(1, 2)$. For any string s, let $\mathbf{v}(s) \in \mathbb{R}^d$ denote its TF–IDF representation under this shared vocabulary and inverse document frequency weighting.

The TF–IDF cosine similarity between a prediction $\hat{y}$ and its reference y is defined as

$$\mathrm{Sim}_{\mathrm{tfidf}}(\hat{y}, y) = \frac{\mathbf{v}(\hat{y})^\top \mathbf{v}(y)}{\|\mathbf{v}(\hat{y})\|_2 \, \|\mathbf{v}(y)\|_2}. \tag{C.6}$$

BERT cosine similarity. We adopt a pretrained sentence embedding model from the Sentence-Transformers framework. Specifically, we use `all-mpnet-base-v2` (Song et al., 2020) to encode each string into a fixed-dimensional vector.

Let $\mathbf{e}(s) \in \mathbb{R}^d$ denote the sentence embedding of a string s produced by the encoder. The BERT cosine similarity between a prediction $\hat{y}$ and its reference y is defined as

$$\mathrm{Sim}_{\mathrm{BERT}}(\hat{y}, y) = \frac{\mathbf{e}(\hat{y})^\top \mathbf{e}(y)}{\|\mathbf{e}(\hat{y})\|_2 \, \|\mathbf{e}(y)\|_2}. \tag{C.7}$$

BERTScore (F1). We use BERTScore as a token-level semantic similarity metric based on contextualized embeddings. Given a prediction $\hat{y}$ and a reference y, both are first tokenized and encoded by a pretrained transformer encoder into contextual token embeddings.

Let $\mathbf{h}_i(\hat{y}) \in \mathbb{R}^d$ denote the embedding of the i-th token in $\hat{y}$, and $\mathbf{h}_j(y) \in \mathbb{R}^d$ the embedding of the j-th token in y. The pairwise token similarity is defined by cosine similarity

$$s_{ij} = \frac{\mathbf{h}_i(\hat{y})^\top \mathbf{h}_j(y)}{\|\mathbf{h}_i(\hat{y})\|_2 \, \|\mathbf{h}_j(y)\|_2}.$$

BERTScore precision and recall are defined as

$$P_{\mathrm{BS}}(\hat{y}, y) = \frac{1}{|\hat{y}|} \sum_i \max_j s_{ij}, \qquad R_{\mathrm{BS}}(\hat{y}, y) = \frac{1}{|y|} \sum_j \max_i s_{ij}. \tag{C.8}$$

The BERTScore F1 is given by

$$\mathrm{BERTScore}_{\mathrm{F1}}(\hat{y}, y) = \frac{2 P_{\mathrm{BS}}(\hat{y}, y) R_{\mathrm{BS}}(\hat{y}, y)}{P_{\mathrm{BS}}(\hat{y}, y) + R_{\mathrm{BS}}(\hat{y}, y)}. \tag{C.9}$$

Appendix B. Automated Evaluation

BLEURT. We adopt BLEURT as a learned neural evaluation model. Formally, let

$$f_{\theta_{\text{BLEURT}}} : \mathcal{U}^* \times \mathcal{U}^* \to \mathbb{R}$$

denote a pretrained neural scoring function that maps a reference action y and a predicted action $\hat{y}$ to a real-valued semantic similarity score.

The BLEURT score is defined as

$$\text{BLEURT}(\hat{y}, y) = f_{\theta_{\text{BLEURT}}}(y, \hat{y}), \tag{C.10}$$

where $f_{\theta_{\text{BLEURT}}}$ is instantiated as a Transformer-based regression model fine-tuned to predict human judgment scores over text pairs. In our experiments, we use the publicly released checkpoint `Elron/bleurt-base-128` hosted on the HuggingFace Model Hub and loaded via the `transformers` library (Wolf et al., 2020). Given a pair $(y, \hat{y})$, the model outputs a single scalar logit, which is used directly as the BLEURT score.

LLM-as-a-Judge score. We implement an LLM-based evaluator to score predicted actions according to a predetermined evaluation rubric.

Formally, let

$$g_\phi : \mathcal{X} \to \mathcal{T}$$

denote a pretrained LLM with fixed parameters ϕ, and let

$$\pi : \mathcal{U}^* \times \mathcal{U}^* \to \mathcal{X}$$

be a deterministic prompting function that maps a prediction–reference pair $(\hat{y}, y)$ into a structured evaluation prompt.

The LLM-as-a-Judge score is defined as

$$\text{Judge}(\hat{y}, y) = \text{Parse}(g_\phi(\pi(\hat{y}, y)))\,, \tag{C.11}$$

where $g_\phi(\pi(\hat{y}, y))$ denotes the model-generated evaluation response, and $\text{Parse}(\cdot)$ extracts a scalar score according to a fixed scoring evaluation rubric. In our experiments, g_ϕ is instantiated as `ChatGPT-4o-mini`, the scoring rubric is defined in Table 2, and all evaluations are performed using deterministic decoding.

Appendix C. Case Study

G Case Study: Repair vs. Replace in Non-Repairable Tire Claims

According to United States. Department of Transportation. National Highway Traffic Safety Administration (2008), punctures located on the tire sidewall are generally considered non-repairable and require tire replacement rather than repair. This rule is well established in automotive safety standards and is routinely enforced in warranty claim operations.

This practice is clearly reflected in our dataset. Among the $2,953$ tire-related claims that contain lexical indicators of sidewall damage (e.g., "inside wall", "in side wall"), only 4 cases resulted repair-type corrective actions (e.g., "repair tire", "patch tire", or "remove nail"), yielding an empirical repair probability of $\frac{4}{2953} \approx 1.36 \times 10^{-3}$. This extreme imbalance closely mirrors real-world repair decisions for sidewall damage. If we assume the repair probability follows a Binomial distribution (not suitable for the imbalance case, again introducing conservative bias), then the expected number of repairs is 4. The fine-tuned model predicts only 1 repair-type actions among the $2,953$ tire-related claims, whereas non-fine-tuned models predict 262 (DeepSeek-R1 + Prompt) and 65 (Qwen-Instruct + Prompt), respectively. These large deviations indicate severe, statistically significant distributional shifts, with non-fine-tuned models systematically overestimating repair probability for practically non-repairable sidewall damage claims.

In contrast, the fine-tuned model's predictions are statistically indistinguishable from the empirical reference distribution. Human evaluation suggests that non-fine-tuned models often rely on generic automotive narratives (e.g., "punctures can often be patched"), reflecting common consumer-level knowledge rather than real-world industry practice. The fine-tuned model, trained on historical warranty claim outcomes, internalizes this highly asymmetric conditional distribution, aligning its prediction with real-world practice.

Appendix D. Research Overview

Cyber Risk

2025 Guo, Zhang, **Quan**. Variance
2025 Chong, Linders, **Quan**, Zhang. ASTIN
🏆 Bob Alting von Geusau Prize

2026 Boyer, **Quan.**

Tree-Based Models

2018 **Quan**, Valdez. DM
2021 **Quan**, Gan, Valdez. AAS
2022 Hu, **Quan**, Chong. IME
2023 **Quan**, Wang, Gan, Valdez. PEIS

2025 Dong, Guillen, Bermudez, **Quan**.

Physics-informed neural networks

2025 Guo, **Quan**.

Climate Risk

InsurTech Innovations

2024 **Quan**, Hu, Dong, Valdez. NAAJ

Automated Machine Learning

2024 Dong, **Quan**. IME

2025 Guo, Dong, **Quan**.
2024 Dong, Liu, **Quan**, Wang.

Natural Language Processing

2024 Zhang, Hu, **Quan**. NAAJ

2024 Dong, **Quan**.
2024 Chen, Cao, **Quan**.
2024 Hu, **Quan**.
2025 Mo, **Quan**.

Federated Learning

2024 Dong, Feng, **Quan**, Wang.