

Article

Artificial Intelligence Adoption and Ethical Governance in Australian Insurance: Evidence from Web-Based Content Analysis

Matias A. Morales Armijo * , Jinhui Zhang  and Yanlin Shi 

Department of Actuarial Studies and Business Analytics, Macquarie Business School, Macquarie University, Sydney, NSW 2109, Australia; colin.zhang@mq.edu.au (J.Z.); yanlin.shi@mq.edu.au (Y.S.)

* Correspondence: matiasalfredo.moralesarmijo@hdr.mq.edu.au

Abstract

The findings of this study indicate that although artificial intelligence (AI) is increasingly embedded in operational practices across the Australian insurance sector, explicit engagement with ethical AI principles remains limited. Based on an analysis of 156 AI-related web pages from Australian insurers, the results show that 58% of companies do not reference any AI Ethics Principles, highlighting a gap between AI adoption and ethical governance. The predominance of operational themes over governance-oriented discourse suggests that ethical considerations are not yet systematically integrated into public-facing communication strategies. While certain principles, particularly human wellbeing and privacy, receive greater attention, others such as accountability and contestability remain comparatively underrepresented. This imbalance indicates a structural gap between technological implementation and transparent ethical articulation. From an actuarial and governance perspective, strengthening the visibility and consistency of responsible AI commitments may enhance stakeholder trust and support sustainable innovation. This study provides one of the first empirical assessments of publicly articulated ethical AI governance in the Australian insurance sector; future research could extend the analysis to international markets and additional data sources.

Keywords: artificial intelligence; insurance; underwriting; Australia; actuarial science; ethics; fairness; transparency; accountability

1. Introduction

Artificial intelligence (AI) has become a transformative force in the insurance industry, reshaping underwriting, claims management, fraud detection and customer engagement. Over the past decade, AI-driven systems have enhanced operational efficiency by automating complex processes and enabling increasingly granular risk assessment. The growing availability of structured and unstructured data from internal systems and external sources, including Internet-of-Things (IoT) devices, allows insurers to refine pricing, personalise policies and support evidence-based decision-making (Eckert and Osterrieder 2020). However, the integration of AI into insurance practice introduces significant governance and ethical challenges that extend beyond technical implementation.

Actuaries play a central role in this transformation. Traditionally responsible for statistical modelling and risk classification, they now increasingly interact with advanced machine learning systems (Zarifis et al. 2019). While AI enhances predictive precision, it also raises concerns regarding algorithmic opacity, bias and accountability. Actuaries therefore operate at the intersection of technological innovation and professional responsibility, where



Academic Editors: Don Hong,
Tianxiang Shi and Chris Groendyke

Received: 21 May 2026

Revised: 1 July 2026

Accepted: 6 July 2026

Published: 20 July 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

ensuring fairness and transparency in AI-supported pricing and underwriting models is essential to maintaining public trust (Richman 2024).

The expansion of AI in insurance has intensified scholarly debate concerning the balance between actuarial fairness and societal fairness. Risk-based pricing seeks to align premiums with individual risk characteristics, yet regulatory frameworks restrict the use of protected attributes such as gender, race or religion. The prohibition of gender-based pricing in the European Union, following Council Directive 2004/113/EC and the Test-Achats ruling, illustrates how actuarial justification may conflict with broader equality objectives (Xin and Huang 2024). Advanced predictive models further complicate this landscape through indirect discrimination, whereby seemingly neutral variables may act as proxies for protected characteristics (Farthing et al. 2021; Holland and Kavuri 2021).

In response, scholars have proposed fairness-aware machine learning and explainable AI (XAI) approaches to mitigate opacity and bias, while regulatory frameworks such as the General Data Protection Regulation (GDPR) reinforce accountability and transparency requirements (Du Preez et al. 2024). Recent studies in the insurance literature have highlighted the growing use of artificial intelligence and machine learning in pricing, underwriting and claims processes, as well as the governance and bias-related risks associated with these applications (Blier-Wong et al. 2021; Amerirad et al. 2023; Mahajan et al. 2025). Nevertheless, despite extensive theoretical and regulatory debate, empirical evidence regarding how insurers publicly articulate and communicate their engagement with ethical AI principles remains limited.

In Australia, the insurance sector has rapidly adopted AI technologies, including generative AI and large language models (Mullins et al. 2021). Although Australia's eight Artificial Intelligence Ethics Principles (Australia's AI Ethics Principles 2025) provide national guidance for responsible AI use, their implementation remains voluntary. As a voluntary national framework, Australia's AI Ethics Principles do not impose sector-specific compliance obligations on insurers; instead, they provide a benchmark for assessing whether organisations publicly signal responsible AI development, deployment, and oversight.

This benchmark is particularly relevant in insurance, where AI-supported underwriting, pricing, claims assessment, and customer engagement can affect fairness, privacy, transparency, contestability, and accountability in decisions that materially influence consumers. The framework supports the responsible development and deployment of AI systems through eight principles: human, societal and environmental wellbeing; human-centred values; fairness; privacy protection and security; reliability and safety; transparency and explainability; contestability; and accountability. Collectively, these principles aim to promote trustworthy AI systems that balance innovation with ethical governance, stakeholder protection, and public trust. Consequently, variation across organisations may increase the risk of inconsistent ethical governance and uneven transparency. Moreover, the way artificial intelligence is discussed may differ across insurance sectors due to differences in business models, data sources and operational priorities. In life and health insurance, the literature frequently emphasises AI applications related to risk assessment, underwriting and the analysis of health or behavioural data, whereas in other insurance contexts, AI is often discussed in relation to pricing, claims management and fraud detection (Eckert and Osterrieder 2020; Boodhun and Jayabalan 2018; Spender et al. 2019). While prior studies examine AI governance conceptually or from regulatory perspectives, empirical evidence on how insurers publicly communicate responsible AI practices remains limited, particularly in the Australian context.

This study addresses this gap by systematically examining how Australian insurance companies articulate their engagement with ethical AI principles in publicly available corporate communications. The research is conceptually structured following the theory–context–characteristics–methodology (TCCM) framework (Hentzen et al. 2022), to structure the analysis of AI adoption and ethical considerations in the Australian insurance industry. Within this framework, *theory* examines the conceptual and ethical foundations guiding AI use, *context* refers to the industry and regulatory environment in which these technologies are implemented, *characteristics* capture the technological features and applications of AI systems, and *methodology* reflects the analytical approaches used to examine their adoption and governance. The empirical assessment relies on web scraping and a mixed-methods approach combining quantitative content analysis, K-means clustering and principal component analysis (PCA). By comparing dominant operational themes with references to ethical governance principles, the analysis evaluates whether technological integration is matched by transparent ethical articulation. The findings contribute empirical evidence to the literature on AI governance and actuarial practice by assessing the current maturity of responsible AI discourse in the Australian insurance sector.

Unlike prior conceptual analyses of AI governance in insurance, the present study offers sector-level empirical measurement of publicly articulated ethical AI commitments.

2. Materials and Methods

2.1. Research Design

This study adopts a mixed methods design to examine the adoption and ethical positioning of artificial intelligence (AI) in the Australian insurance industry, with particular attention to actuarial and underwriting contexts. The analysis integrates quantitative content frequency measures with qualitative thematic assessment, guided by the Theory–Context–Characteristics–Methodology (TCCM) framework (Hentzen et al. 2022).

The methodological design was structured to maximise transparency, replicability, and minimise subjective bias in principle identification and thematic interpretation.

2.2. Population and Search Scope

The population comprised all Australian insurers listed by regulatory and industry bodies at the time of data collection (25 March 2025), including:

- 52 general insurers (Insurance Council of Australia)
- 34 life insurers (APRA Register of Life Insurance Companies)
- 30 private health insurers (APRA Register of Private Health Insurers)

Only Australian domains (.au) were considered.

Inclusion Criteria

- Media releases and official web pages
- Explicit reference to AI adoption or AI-related initiatives
- English language content

Exclusion Criteria

- Non-Australian sources
- PDFs or non-web-native documents in the initial search phase
- Websites preventing automated access due to security restrictions

2.3. Data Collection Procedure

Data collection followed a two-phase automated protocol.

Phase 1: Web Search and Filtering

Search queries combined company names with the terms:

“insurance” AND “artificial intelligence” AND (“news” OR “breaking” OR “report”)

Google Cloud API and Google Programmable Search Engine were used to retrieve URLs. The initial output (“founded_URL”) was filtered programmatically to retain only pages containing both “insurance” and “artificial intelligence”, producing a refined dataset (“successful_URL”).

In total:

- 242 URLs were initially retrieved
- 156 URLs met filtering criteria

The number of retrieved URLs varied across companies depending on the availability of publicly accessible AI-related content. Consequently, insurers with a larger digital presence may have contributed a greater number of webpages to the final dataset. The study focuses on publicly available AI-related communications rather than attempting to standardise the number of webpages per company, and this should be considered when interpreting the results.

Phase 2: AI Principle Identification

The presence of Australia’s 8 AI Ethics Principles (Australia’s AI Ethics Principles) was examined across retrieved documents.

A structured keyword-based mapping system was developed, consisting of 16 predefined keywords (two per principle) organised as semantically related pairs to capture lexical variation while preserving conceptual alignment. Each principle was ultimately represented by one selected keyword for final analysis. The selected keywords were intended to function as representative indicators of each principle rather than exhaustive dictionaries of all possible expressions. This approach was adopted to maintain consistency and replicability across the analysis. Nevertheless, some implicit or alternative formulations of ethical concepts may not have been captured through the selected terminology.

Further details on the keyword selection and validation process, including inter-coder agreement, are provided in Section 2.4.

The predefined keyword pairs used for mapping each principle are presented in Table 1.

Table 1. Keyword pairs used to map AI principles.

Principle	Keyword Pair
Human, societal and environmental wellbeing	Benefit/wellbeing
Human-centred values	Rights/autonomy
Fairness	Inclusivity/equity
Privacy protection and security	Privacy/security
Reliability and safety	Reliability/robustness
Transparency and explainability	Transparency/explainability
Contestability	Challenge/appeal
Accountability	Accountability/responsibility

2.4. Inter-Coder Reliability

To ensure methodological rigor, two independent coders (the researcher and an academic expert) selected one representative keyword per principle from the predefined pairs. The selections were made independently to ensure objectivity.

The selected keywords were then mapped to their corresponding Australian AI Ethics Principle categories using a predefined mapping dictionary developed by the authors. The dictionary was constructed by manually assigning representative keywords and synonymous terms to each of the eight Australian AI Ethics Principles based on the official descriptions of the Australian AI Ethics Principles published by the Australian Government. This mapping was established prior to the analysis and applied consistently across all webpages to ensure systematic and reproducible keyword classification. For the purpose of agreement analysis, each principle category was encoded numerically and treated as nominal data. Inter-coder agreement was assessed using Krippendorff's alpha (Krippendorff 2011). The resulting coefficient was $\alpha = 0.743$.

This indicates substantial agreement and supports the reliability of the keyword classification framework.

2.5. Text Processing and Vector Database Construction

For each URL, web content was programmatically extracted and cleaned. Text was segmented into fixed-length units to facilitate structured retrieval and indexed within a vector database environment. The computational pipeline was implemented in Python 3.11 using open-source libraries for text processing and semantic indexing.

2.6. Large Language Model-Assisted Retrieval

A retrieval-based approach was used to detect predefined keywords and their direct synonyms within document contexts.

An open-source large language model (Llama 3.1–8B) (Grattafiori et al. 2024) was deployed locally to assist structured retrieval. The model operated under predefined prompt constraints designed to detect exact keyword matches or direct synonyms only. It was explicitly restricted from performing open-ended semantic expansion or autonomous classification. The model served solely as a controlled retrieval and verification tool and did not generate synthetic content or alter source material.

2.7. Quantitative and Qualitative Analysis

Quantitative Analysis

Frequency counts of principle-related keywords were computed to assess relative prominence across insurer types (general, life, health). Sentiment classification was performed using VADER (Hutto and Gilbert 2014). Although VADER was originally developed for social media text, its rule-based lexicon approach allows consistent detection of general valence patterns in formal corporate communications.

Qualitative Analysis

Thematic analysis was applied to contextual excerpts retrieved by the model to evaluate how principles were framed and operationalised.

2.8. Limitations

Some company websites could not be accessed due to automated bot protections or security configurations. Additionally, the keyword-based approach relied on representative terms associated with each principle and may not capture all linguistic variations or implicit ethical commitments expressed using alternative terminology.

Furthermore, the clustering solution produced a relatively low Silhouette Score, indicating overlap between thematic groups. Consequently, the identified clusters should be interpreted as exploratory thematic structures rather than definitive classifications.

3. Results

3.1. Data Extraction Analysis

From the initial population of 116 Australian insurers (52 general insurers, 34 life insurers and 30 health insurers), a total of 156 AI-related webpages were retrieved and met the inclusion criteria described in Section 2.3. These webpages constituted the analytical dataset used for thematic clustering and sentiment analysis. The textual content extracted from these webpages was then analysed to determine whether it could be grouped into coherent thematic clusters. K-means clustering was applied because it is a widely used unsupervised learning technique for identifying latent structures in large textual datasets and grouping similar documents based on semantic similarity. The optimal number of clusters (k) was evaluated using both the elbow method and the Silhouette Score (Géron 2023). To maintain interpretability and enable meaningful manual interpretation of the resulting themes, the search range was restricted by imposing a maximum value of $k = 20$. This constraint also avoids excessive fragmentation of topics, which may reduce the analytical clarity of thematic interpretations.

The elbow method did not show a clear inflection point. However, the Silhouette Score reached its maximum at $k = 16$ (Figure 1), indicating $k = 16$ as the optimal configuration according to the Silhouette criterion. Although $k = 16$ achieved the highest Silhouette Score among the evaluated solutions, the absolute value of 0.14 remains relatively low. This suggests limited separation between clusters and indicates that the resulting themes should be interpreted as exploratory patterns rather than strongly distinct thematic categories. The Silhouette values for selected k are summarised in Table 2.




Figure 1. Elbow Method and Silhouette Score for different values of k .

Table 2. Silhouette Score comparison across selected values of k .

K	Silhouette Score	Visual Spread	Cluster Cohesion	Interpretability
5	0.059	Poor	Overmerged	Very easy
10	0.094	Decent	Slight overlap	Easy
15	0.115	Good	More structure	Still doable
16	0.14	Best so far	Clearer regions	Slightly harder
17	0.122	Moderate	Slightly noisier	Getting complex
20	0.133	Fragmented	Mixed cohesion	Hard to label

Because the elbow method did not identify a clear inflection point, the Silhouette Score was adopted as the primary quantitative criterion for selecting the clustering solution. To validate this selection, PCA visualisations were produced for $k = 5, 10, 15, 16, 17$, and 20

(Figure 2). The projection corresponding to $k = 16$ demonstrated comparatively greater visual separation and cluster cohesion. Accordingly, $k = 16$ was selected as the final clustering solution based on the highest Silhouette Score, together with a qualitative assessment of cluster separability.




Figure 2. PCA cluster comparison across selected k values.

3.1.1. Cluster Labelling and Thematic Aggregation

Each of the 16 clusters was manually labelled based on the 30 most frequent terms, excluding company names (see Section 2.4 for reliability assessment). For example, Cluster 5 contained terms such as *information*, *data*, *personal*, *privacy*, *breach*, and *fraud*, and was labelled Data Privacy and Customer Security.

The complete list of clusters and associated themes is summarised in Table 3.


To improve interpretability, the 16 clusters were aggregated into six higher-level themes based on semantic similarity and domain alignment:

1. Finance and Risk
2. Insurance Services
3. Health and Wellbeing
4. Industry and Technology
5. Data and Security
6. Government and Social Services

The distribution of documents across themes is shown in Figure 3. Finance and Risk accounts for 50.7% of all documents, while Data and Security and Government and Social Services each represent less than 5%.

Table 3. Cluster labels and thematic grouping.

Name Cluster	n_URLs	n_words	Theme
Social Media and Digital Communication	4	5988	Data and Security
Actuarial Science and Risk	16	152,891	Finance and Risk
Healthcare Costs and Informed Consent	6	13,076	Health and Wellbeing
Audit and Financial Oversight	25	290,143	Finance and Risk
Funeral Services and Planning	9	24,660	Insurance Services
Data Privacy and Customer Security	3	11,979	Data and Security
Injury and Claims Management	9	32,507	Insurance Services
Investments and Financial Markets	36	534,683	Finance and Risk
Digital Tools and Innovation in Insurance	6	8233	Industry and Technology
Mental Health and Life Insurance	3	4471	Health and Wellbeing
Personal and Lifestyle Insurance Products	12	23,743	Insurance Services
Motor Vehicles and Engineering	5	14,038	Industry and Technology
Veteran and Defence Services	2	1874	Government and Social Services
Medical Practice and Professionals	10	14,970	Health and Wellbeing
Healthcare and Consent Practices	4	17,873	Health and Wellbeing
Government Procurement and Insurance Markets	2	8318	Government and Social Services

**Figure 3.** Distribution of documents by thematic cluster.

The word cloud for the Finance and Risk theme (Figure 4) highlights dominant terms such as *financial*, *investment*, *super*, *equity*, and *ASX*, highlighting frequent references to financial markets and investment-related terminology.

The PCA projection in Figure 5 provides a visual representation of relative similarity patterns between clusters. As PCA is a dimensionality reduction technique, distances in the two-dimensional projection should be interpreted cautiously as indications of relative similarity rather than measures of exact semantic distance. Finance-related clusters appear more dispersed, suggesting greater variability and more specialised terminology. In contrast, clusters associated with Insurance Services, Health and Wellbeing, and Data and Security are more tightly grouped, indicating higher similarity in their underlying vocabulary and thematic content. Government and Social Services and Industry and Technology also appear relatively close to this central grouping, suggesting moderate overlap with other themes.



Figure 4. Word cloud for the Finance and Risk theme.




Figure 5. PCA projection of thematic clusters.

3.1.2. Sentiment Analysis

Sentiment analysis was used to assess the tone of insurers' public communications about AI-related initiatives and governance issues. Sentiment was classified using VADER (Hutto and Gilbert 2014), a rule-based tool that assigns polarity scores based on a predefined lexicon and linguistic heuristics. Because VADER was originally developed for social media text, the results are interpreted as exploratory indicators of communication tone rather than as evidence of insurers' internal attitudes or governance maturity.

Figure 6 presents the distribution of positive, neutral, and negative sentiment across general, health, and life insurers.

Overall, 83.2% of content was classified as positive, 14.8% as negative, and 1.9% as neutral. Life insurers accounted for the largest proportion of positively classified content (39.4%).

The most frequent terms in negatively classified texts are shown in Figure 7. Words such as claim, cost, and compensation appear prominently within the negatively classified subset. Negative sentiment classifications should be interpreted cautiously, as

Table 4. AI Ethics Principle identification by insurance category.

Principle References	General	Health	Life	Total	% of Total Observations
No Principle	55	34	77	166	58%
Principle References	41	30	51	122	42%
Grand Total	96	64	128	288	100%

Table 5. Frequency of principle mentions by company type.

Principles	General	Health	Life	Total	% Total Principle/Total
Human, societal and environmental wellbeing	10	8	13	31	25%
Privacy protection and security	7	6	14	27	22%
Human-centred values	6	4	7	17	14%
Contestability	5	4	6	15	12%
Reliability and safety	4	3	3	10	8%
Transparency and explainability	3	3	3	9	7%
Accountability	3	1	3	7	6%
Fairness	3	1	2	6	5%
Grand Total	41	30	51	122	42%

The distribution of principle references at the company level reveals variation across companies within each insurance sector. While comprehensive coverage across all eight principles is observed in a small number of cases, most insurers exhibit considerable variation in the public disclosure of AI Ethics Principles, reflecting differences in the extent to which organisations publicly disclose their AI governance practices (Tables 6–8). The Australian AI Ethics Principles are voluntary guidelines rather than legally binding requirements. Therefore, the absence of references to these principles in publicly available company documents should not be interpreted as evidence of regulatory non-compliance. Instead, the findings reflect differences in the extent to which insurers publicly disclose their AI governance practices. Although not legally required, publicly disclosing AI governance practices may enhance organisational transparency, stakeholder trust, and accountability. Conversely, limited public disclosure may reduce transparency regarding how organisations disclose their approach to responsible AI.

A comparison between the number of AI-related web pages and principle mentions (Figure 8) indicates no consistent relationship between content volume and ethical principle coverage, suggesting that increased AI communication efforts do not necessarily reflect stronger engagement with ethical principles.

Table 6. Principles Identified for General Insurance Companies.


Company	Principles
Auto & General Insurance Company Limited	8
HDI Global Specialty SE	5
Hollard Insurance Partners Limited	5
ANZ Lenders Mortgage Insurance	4
Allianz Australia Insurance Limited	3
Guild Insurance Limited	3
Insurance Australia Group Limited	3
Pacific International Insurance Pty Ltd.	3
AAI Limited	2
QBE Insurance (Australia) Limited	2
Zurich Financial Services Australia Limited	2
RACQ Insurance Limited	1

Table 7. Principles Identified for Health Insurance Companies.

Company	Principles
Medibank Private Limited	8
Defence Health Limited	6
CBHS Health Fund Limited	5
HBF Health Limited	3
Queensland Teachers' Union Health Fund Limited	3
Australian Unity Health Limited	2
BUPA HI Pty Ltd.	2
Westfund Limited	1

Table 8. Principles Identified for Life Insurance Companies.

Company	Principles
Hannover Life Re of Australasia	7
NobleOak Life	7
MLC Limited	6
Resolution Life Australasia	5
AIA Australia	3
Combined Life Insurance Company of Australia	3
Futurity Investment Group	3
ClearView Life Assurance	2
Foresters Financial	2
HCF Life Insurance Company Pty	2
Lifeplan Australia Friendly Society	2
Munich Reinsurance Company of Australasia	2
RGA Reinsurance Company of Australia	2
Swiss Re Life & Health Australia	2
Zurich Australia	2
Allianz Australia Life Insurance	1

**Figure 8.** Number of AI-related web pages and corresponding principle mentions by company.

Sectoral Comparison

Principle adoption patterns across insurance categories are illustrated in the radar chart (Figure 9) and grouped bar chart (Figure 10). Health insurers demonstrate comparatively higher representation across multiple principles, while life insurers show a stronger emphasis on privacy-related references. General insurers display a more even distribution across principles, although at lower overall frequencies.




Figure 9. Radar chart of principle adoption.




Figure 10. Adoption of ethical principles by insurance type.

Privacy Protection and Security is particularly prominent in life (88%) and health (75%) insurance. In contrast, Fairness and Accountability remain below 25% across

all sectors, indicating comparatively limited public articulation of distributive and governance-oriented AI safeguards.

Collectively, these findings demonstrate thematic concentration in operational domains and variation in the frequency of ethical AI principle references across the Australian insurance sector.

4. Discussion

The findings reveal a structural asymmetry in how Australian insurers publicly communicate about artificial intelligence. Thematic clustering demonstrates that AI discourse is predominantly concentrated on operational and financially oriented applications, particularly within finance and risk, insurance services, and health-related domains. This pattern is consistent with prior research suggesting that AI adoption in insurance has been primarily driven by efficiency gains and automation (Zarifis et al. 2019; Singh and Chivukula 2020). In contrast, governance-oriented themes appear comparatively less visible, indicating that public AI communication remains largely application-focused rather than explicitly governance-driven.

The analysis of ethical principle adoption reinforces this observation. A majority of insurers did not reference any of the eight Australian AI Ethics Principles, and among those that did, emphasis was concentrated on wellbeing and privacy, principles closely aligned with established regulatory obligations. Fairness and accountability, however, were comparatively underrepresented. Given the increasing role of AI in underwriting, pricing and claims assessment, the limited public articulation of these principles suggests uneven maturity in the external communication of responsible AI practices.

Sectoral variation provides additional nuance. Health insurers demonstrated comparatively broader engagement with human-centred principles, while life insurers emphasised privacy and data protection. General insurers exhibited a moderate but more uniform distribution across principles. These differences indicate that ethical AI articulation may reflect sector-specific risk exposures and regulatory environments rather than a uniform industry trajectory.

Importantly, the absence of a consistent relationship between the volume of AI-related content and ethical principle references suggests that communication intensity does not necessarily correspond to governance transparency. These findings suggest that public ethical articulation may reflect disclosure strategies rather than directly mirroring internal governance maturity.

The predominance of positive sentiment suggests that AI is framed primarily as innovation-enhancing. However, given the regulatory and ethical complexity of AI deployment in underwriting and claims assessment, this positivity may reflect strategic communication rather than comprehensive governance transparency. Prior scholarship emphasises that explainability and accountability must be embedded throughout the AI lifecycle to sustain trust and contestability (Mullins et al. 2021; Owens et al. 2022). The comparatively limited emphasis on these principles suggests that responsible AI integration remains an evolving domain within Australian insurance.

Several limitations warrant consideration. Keyword-based identification may not capture implicit ethical commitments, and sentiment models may misclassify sector-specific terminology. Moreover, web-based communications reflect public discourse rather than internal governance mechanisms. Future research could incorporate longitudinal analysis, semantic modelling techniques, and triangulation with regulatory disclosures or interviews to provide deeper insight into the operationalisation of responsible AI practices.

Overall, this study contributes empirical evidence to the literature on AI governance by demonstrating that, within the Australian insurance sector, operational AI integration

appears more visible than explicit ethical articulation. Strengthening governance mechanisms and enhancing transparent communication may support stakeholder trust and alignment with evolving regulatory expectations. Importantly, the analysis evaluates publicly articulated governance signals rather than internal compliance structures or operational AI controls.

5. Conclusions

This study demonstrates that, within Australian insurance, public AI discourse remains predominantly operational, while explicit ethical articulation is comparatively limited and uneven. While AI adoption appears operationally advanced across life, health and general insurance sectors, explicit engagement with ethical governance principles is less consistently demonstrated.

Empirical analysis indicates that AI-related content is concentrated on financial and service-oriented themes, with limited and selective reference to Australia's eight AI Ethics Principles. Although 42% of analyzed insurers mentioned at least one principle, engagement was concentrated primarily on wellbeing and privacy, whereas fairness, accountability, transparency and explainability were comparatively underrepresented.

From a risk governance perspective, this asymmetry between technological integration and ethical articulation may present emerging reputational and regulatory challenges as expectations for responsible AI oversight continue to evolve. By offering a replicable methodological framework for assessing ethical AI communication, this research contributes to the growing literature on AI governance within regulated financial industries.

Future research should extend beyond corporate web discourse to incorporate longitudinal analysis, stakeholder perspectives and international comparisons, thereby deepening understanding of how responsible AI practices are operationalised and communicated in practice.

Author Contributions: Conceptualization, M.A.M.A. and J.Z.; Methodology, M.A.M.A.; Validation, M.A.M.A.; Formal analysis, M.A.M.A.; Investigation, M.A.M.A.; Resources, J.Z. and Y.S.; Writing—original draft, M.A.M.A.; Writing—review & editing, M.A.M.A., J.Z. and Y.S.; Visualization, M.A.M.A.; Supervision, J.Z. and Y.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The dataset generated and analysed during the current study is not publicly available due to third-party website terms of use, copyright, licensing, and data redistribution restrictions associated with the web-scraped URLs and keyword mappings used in this study.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Amerirad, Behnaz, Matteo Cattaneo, Ron S. Kenett, and Elisa Luciano. 2023. Adversarial Artificial Intelligence in Insurance: From an Example to Some Potential Remedies. *Risks* 11: 20. [CrossRef]
- Australia's AI Ethics Principles. 2025. Department of Industry, Science and Resources. Australia's AI Ethics Principles. Available online: <https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles> (accessed on 1 April 2024).
- Blier-Wong, Christopher, Hélène Cossette, Luc Lamontagne, and Etienne Marceau. 2021. Machine Learning in P&C Insurance: A Review for Pricing and Reserving. *Risks* 9: 4.
- Boodhun, Noorhannah, and Manoj Jayabalan. 2018. Risk prediction in life insurance industry using supervised learning algorithms. *Complex & Intelligent Systems* 4: 145–54. [CrossRef]

- Du Preez, Valerie, Shaun Bennet, Matthew Byrne, Aureli  n Couloumy, Arijit Das, Jean Dessain, Richard Galbraith, Paul King, Victor Mutanga, Frank Schiller, and et al. 2024. From bias to black boxes: Understanding and managing the risks of AI—An actuarial perspective. *British Actuarial Journal* 29: e6. [\[CrossRef\]](#)
- Eckert, Christian, and Katrin Osterrieder. 2020. How digitalization affects insurance companies: Overview and use cases of digital technologies. *Zeitschrift f  r die Gesamte Versicherungswissenschaft* 109: 333–60. [\[CrossRef\]](#)
- Farthing, Sophie, John Howell, Katerina Lecchi, Zoe Paleologos, Phoebe Saintilan, and Edward Santow. 2021. *Human Rights and Technology Final Report*. Geneva: Human Rights Commission.
- G  ron, Aur  lien. 2023. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow THIRD EDITION Concepts, Tools, and Techniques to Build Intelligent Systems*. Sebastopol: O'Reilly Media, Inc.
- Grattafiori, Aaron, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and et al. 2024. The Llama 3 Herd of Models. *arXiv* arXiv:2407.21783.
- Hentzen, Janin Karoli, Arvid Hoffmann, Rebecca Dolan, and Erol Pala. 2022. Artificial intelligence in customer-facing financial services: A systematic literature review and agenda for future research. *International Journal of Bank Marketing* 40: 1299–336. [\[CrossRef\]](#)
- Holland, Christopher P., and Anil Kavuri. 2021. *Artificial Intelligence and Digital Transformation of Insurance Markets*. Cham: Palgrave Macmillan.
- Hutto, Clayton, and Eric Gilbert. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. In *Proceedings of the International AAAI Conference on Web and Social Media*. Washington: AAAI Press, pp. 216–25.
- Krippendorff, Klaus. 2011. *Computing Krippendorff's Alpha-Reliability*. Thousand Oaks: Sage.
- Mahajan, Siddharth, Rohan Agarwal, and Mihir Gupta. 2025. Algorithmic Bias Under the EU AI Act: Compliance Risk, Capital Strain, and Pricing Distortions in Life and Health Insurance Underwriting. *Risks* 13: 160. [\[CrossRef\]](#)
- Mullins, Martin, Christopher P. Holland, and Martin Cunneen. 2021. Creating ethics guidelines for artificial intelligence and big data analytics customers: The case of the consumer European insurance market. *Patterns* 2: 100362. [\[CrossRef\]](#) [\[PubMed\]](#)
- Owens, Emer, Barry Sheehan, Martin Mullins, Martin Cunneen, Juliane Ressel, and German Castignani. 2022. Explainable Artificial Intelligence (XAI) in Insurance. *Risks* 10: 230. [\[CrossRef\]](#)
- Richman, Ronald. 2024. *An AI Vision for the Actuarial Profession*. Arlington: CAS Forum.
- Singh, Sushant K., and Muralidhar Chivukula. 2020. A Commentary on the Application of Artificial Intelligence in the Insurance Industry. *Trends in Artificial Intelligence* 4: 75–79. [\[CrossRef\]](#) [\[PubMed\]](#)
- Spendler, Anna, Colin Bullen, L. Altmann-Richer, James Cripps, Robin Duffy, Chris Falkous, Mark Farrell, Tony Horn, James Wigzell, and Wendy Yeap. 2019. Wearables and the internet of things: Considerations for the life and health insurance industry. *British Actuarial Journal* 24: e22. [\[CrossRef\]](#)
- Xin, Xi, and Fei Huang. 2024. Antidiscrimination Insurance Pricing: Regulations, Fairness Criteria, and Models. *North American Actuarial Journal* 28: 285–319. [\[CrossRef\]](#)
- Zarifis, Alex, Christopher P. Holland, and Alistair Milne. 2019. Evaluating the impact of AI on insurance: The four emerging AI- and data-driven business models. *Emerald Open Research* 1: 15. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.