AI INDUSTRY INSIGHTS:

Business Value and Governance Strategies

LIMRA LOMA

Navigate With Confidence

Contents

1.0 Overview

THE FINANCIAL SERVICES INDUSTRY is facing a profound transformation fueled by advancements in artificial intelligence (AI), including a key subset of the broader field — generative AI (GenAI). Like other sectors, the rapid growth of AI has drawn significant attention, prompting companies to invest time and resources in harnessing its potential. In the insurance industry, firms are striving to grasp AI's capabilities and determine how to leverage it for long-term success.

AI stands ready to revolutionize various aspects of the insurance value chain. The integration of AI and GenAI in the life insurance industry is as much about reengineering organizational business processes to take advantage of the technology as it is about the technology itself. The introduction of AI across the value chain is not just enhancing existing processes but is fundamentally reshaping the life insurance landscape. When leveraged effectively, AI can provide insurers with substantial competitive advantages, foster innovation, and enhance their ability to meet the needs of a new generation of customers in the digital age.

This study is focused on examining the current state of AI business value enablement (use cases) and AI governance, offering a snapshot of AI adoption and implementation across the industry in the first half of 2024. As the financial services industry seeks to define excellence, understanding the present landscape is critical — from best practices in AI deployment to risk management and value measurement. As the first comprehensive study of its kind in our sector, it not only provides a clear picture of where the industry stands today, but also charts a course for future advancements. These findings will help guide the development of best practices, frameworks, and tools to support the industry's journey from its current state to its desired future.

Geared towards AI leaders (business and technology) and practitioners, this report is complementary to two concurrently released reports: *Navigating the AI Landscape: Current State of the Industry — Executive Briefing* and *The AI Industry Today: Understanding the Current State of Play*.

Kartik Sakthivel, Ph.D., MS-IT/MS-CS, MBA, PGC-IQ
Vice President & Chief Information Officer, LIMRA and LOMA
ksakthivel@limra.com

2.0 AIGG Phase One Approach and Methodology

THE PHASE ONE CURRENT STATE ASSESSMENT STUDY is a culmination of a questionnaire issued to members of the LIMRA and LOMA AI Governance Group (AIGG), individual one-on-one conversations with business and technology executives within and outside the AIGG, as well as by dialogue with consulting organizations, ecosystem partners, and analyst firms. Data collected through these means has been compiled with a special emphasis to ensure that any recommendations rendered are insightful and actionable over the short and medium terms. Over two dozen articles, whitepapers, and research reports published by consulting firms (including the "Big Four") were reviewed to ensure that any redundancy in provided information was kept to a minimum and that the insights being delivered are distinct, unique, and germane to our industry.

The questionnaire (see below), developed by LIMRA and LOMA, was posed to the AIGG members, and served as the central predicate of the current state assessment.

FOCUS AREA: BUSINESS VALUE ENABLEMENT (USE CASES)

1. What are the use cases currently being leveraged in your firm (published information will be sanitized and anonymized)?
2. Use cases currently experiencing success in your firm (published information will be sanitized and anonymized and/or categorized).
3. Measuring AI success — cost-benefit analyses of use cases and ROI measurements. How are you measuring value?
4. Is your organization building AI, buying/renting AI, or both?
5. What kind of processes are you using to shepherd successful POCs/pilots through operationalization?
6. What happens when an experiment is successful? How does it scale?
7. How are AI-related enablers being featured in your corporate goals and objectives?
8. How has the explosion of AI, and GenAI specifically, altered your 1 to 3 year technology roadmaps and/or investments?

FOCUS AREA: GOVERNANCE

1. What is the current governance model within your firm (4Ps: policies, protocols, procedures, processes)?
2. Have you built an AI risk profile that accounts for cybersecurity, data privacy and protection, regulatory, and compliance? How is your firm preparing for AI regulation?
3. What is the state of employee education on AI? Are you leveraging an AI literacy program?
4. What is the state of data readiness for AI within your firm? Do you have a data strategy and governance program? How is your data literacy?
5. How are you planning to manage the cultural change curve with adoption of AI within your firm?
6. How are you governing and derisking the insurance supply chain? That is, how are you governing suppliers and vendors that are leveraging AI within their products and services?
7. What is your organizational thinking around ensuring that AI is transparent and explainable?
8. What are some of the ways your firm is going about ensuring that AI is free of bias and/or proxy discrimination?

3.0 Key Takeaways

PHASE ONE of the LIMRA and LOMA cross-industry AI Governance Group's work was focused on conducting a current state assessment. Each of the key takeaways listed here includes a segment called ***Implications***, which seeks to infer what a specific key takeaway means to the current state of AI within the industry. Additionally, each key takeaway also includes a segment called ***Actionable Insights***. This is intended to present business and technology leaders with specific actions to consider planning and executing within the next 18 to 24 months, commensurate to each key takeaway.

Section 3.1 explores which AI — particularly GenAI, use case domains are realizing value in the industry today. Section 3.2 explores how firms are operating in the absence of overarching regulation or regulatory frameworks to base their AI governance policies on. A supplementary report issued concurrent to this report, *The AI Industry Today: Understanding the Current State of Play*, presents a holistic look at the AI state of play across the entire industry as of mid-2024.

3.1 Business Value Enablement (Use Cases)

Nearly 100 percent of carriers as of Q2 2024 are experimenting with AI/GenAI use cases at varying levels, with about 75 percent piloting and/or planning to pilot or implementing internal GenAI capabilities in 2024. Most of them have established internal cross-company consortia/centers of excellence that are focused on experimenting with GenAI use cases. Some firms are deriving value from their first-generation use cases, while concurrently developing their enterprise AI strategies. Firms that have established and rapidly matured their AI cross-company teams have also sought to implement governance structures around AI within their organizations.

GenAI has been a catalyst that has allowed firms to centralize AI use-case development efforts. Although most firms had already been pursuing AI use cases, they had been limited to specific domains (such as underwriting, marketing, actuarial, etc.), and as such, embedded within individual business units. Although low to moderate coordination at an enterprise-level existed to gain modest efficiencies across these siloed use cases, it had not been anywhere at the level of centralization and coordination as of Q2 2024. This has benefited areas ripe for business process reengineering and/or can derive value from GenAI. Centralized coordination of the use-case development pipeline is allowing firms to experiment with GenAI in parts of the firm that would have otherwise been unlikely to experience the impact of AI.

Leaders should be clear about the fact that while GenAI can be invaluable for a lot of use cases, it will not add value for every use case. Companies are exploring a much wider range of AI/GenAI applications, including fraud detection (through pattern recognition and anomaly detection), claims automation (by automating various manually rigorous aspects of the claims process), customer service enhancements (via chatbots and virtual assistants that handle basic customer questions, provide policy information, assist with claims, etc.), underwriting prediction models (including automated/accelerated

RAPID READ:

As of Q2 2024, most carriers are experimenting with AI use cases, focusing on deriving value from first-generation use cases while developing enterprise AI strategies. The centralization of AI use case development efforts for efficiency and broader applicability is going to be vital. Firms need to develop a clear cost-benefit analysis strategy. It is highly likely that most companies will arrive at a hybrid approach in the "build vs. buy" decision for AI models. As AI programs mature, it will be critical for carriers to align their AI strategies with broader corporate objectives. Agility and adaptability in the industry will be more important now than ever before. Our industry should not focus only on keeping humans in the center of all AI decisions for the near future, but also keep sight of core tenets of humanity as the center of customer-facing AI implementations.

underwriting and risk assessment by development of individual risk profiles), and operational efficiency improvements. There's a significant emphasis on leveraging AI to enhance decision-making, reduce fraud, and improve customer interactions. In discussions with carriers, more than 8 in 10 of large carriers state that productivity gains are their primary driver for GenAI implementation, with nearly two thirds of all those with ongoing pilots expecting cost savings as a direct benefit.

Carriers are realizing successes from their AI use cases. They vary from productivity savings, reduced turnaround times, enhanced customer experiences, cybersecurity improvements, and transcribing meeting summaries, to legacy code modernization, data mapping and migration, and the ability to assist IT developers to author more accurate code, faster. Several companies are in the early stages of AI adoption, with a mix of experimental and operational use cases. Whereas carriers have (in some instances) upwards of 200 use cases in their experimentation pipeline, a majority of carriers are placing priority on AI use cases that are transformative to certain domains across the value chain. Reflective of the fact that carriers had ongoing AI use cases embedded within domains that have since been made part of the enterprise AI consortia, these use cases focus on transforming distribution, claims, and underwriting, etc. There are opportunities for reuse across what used to be siloed AI use cases and some carriers are trying to "build once, use everywhere." GenAI use cases are being prioritized based on where the greatest opportunities are to yield fastest results and broadest cross-enterprise applicability.

Carriers are challenged to effectively measure the success of their AI use cases pipeline. At least in the near term, this will continue to be a challenge. Several factors are serving as headwinds to crafting and executing an effective cost-benefit analysis strategy to measure returns on AI investments. First, most firms have never specifically tracked AI as an investment item within their technology budgets. AI has been incorporated as part of technology budgets, or to complicate things further, been a part of business unit expense line items (for example, underwriting and marketing might have maintained budgets for two different vendor platforms, both delivering services with AI as the enabling technology). Second, with multiple commonly used technology vendors/providers now incorporating AI into their platforms, some at modest licensing costs, carriers will need to understand how to delineate what they are already paying

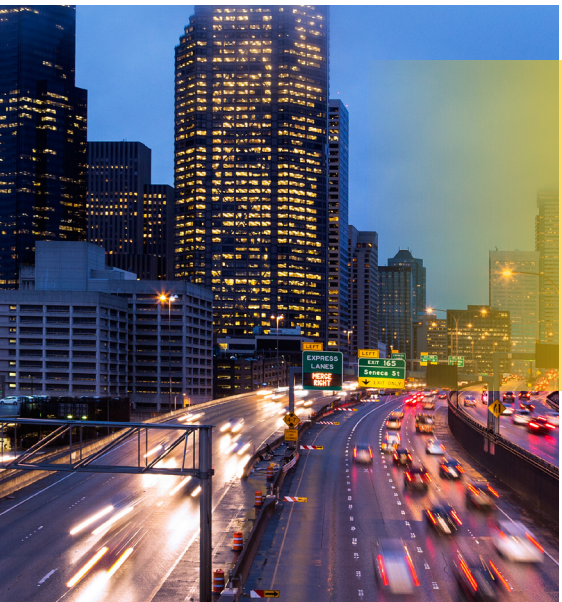

More than* 8 in 10 *large carriers state that productivity gains are their primary driver for GenAI implementation, with nearly two thirds of all those with ongoing pilots expecting cost savings as a direct benefit.

for licensing or subscribing to these vendor platforms versus what they would be investing specifically for the AI add-ons. To measure value for AI that is fully embedded in existing platforms would be implausible. Carriers are actively devising AI value measurement strategies, drawing from traditional means to measure derived value of technology. These strategies include cost-benefit analyses (CBAs), productivity gains tracking, and accuracy improvements in decision-making processes. Some companies have yet to establish formal value measurement frameworks, suggesting that the industry is still exploring how best to quantify AI's impact.

A "build vs. buy" decision presents carriers with the fundamental choice: build out AI models in-house and use their own data for these models along with data sets purchased from trusted external providers, or buy services from technology providers/third-party vendors, wherein the vendors build, manage, and maintain the AI models, as well as manage the data that fuels these AI models. Companies are settling on a hybrid approach in terms of a "build vs. buy" decision, buying large language models (LLMs), while building specific applications on top of them. Staying with the hybrid theme, companies are using various approaches in their AI implementation processes, from proof of concept (POC) and pilot testing to scaling successful initiatives. These varied approaches include test and learn processes, oversight by an AI task force (also known as the Center of Excellence and/or Governance Group), and prioritization based on business impact. More than half of carriers report that they have been diligent in ensuring that their AI strategies often align with broader corporate objectives. These goals are sometimes directly tied to the CEO's focus or the organization's technological roadmap. The explosion of AI —and GenAI — seems to have altered carriers' 1-to-3-year technology roadmaps and/or investments. While there is insufficient data to quantify the full impact of budgets being redirected to AI (an opportunity for LIMRA and LOMA to establish benchmarks), given that there is heightened awareness and interest for AI/GenAI use cases at the CEO/Board level, some carriers report that there have been specific budget allocations for AI exploration. Several report that they have effectuated significant shifts in their technology strategy to accommodate AI initiatives.

A "build vs. buy" decision presents carriers with the fundamental choice: build out AI models in-house and use their own data for these models along with data sets purchased from trusted external providers, or buy services from technology providers/third-party vendors, wherein the vendors build, manage, and maintain the AI models, as well as manage the data that fuels these AI models.

IMPLICATIONS

Enterprise Agility and Adaptability Takes Center Stage: Ours is a conservative, traditionally risk averse industry. It is therefore remarkable to witness how expeditious carriers have been with AI use case experimentation. For organizations that have not been "early adopters" or "fast followers," the so-called FOMO (fear of missing out) effect is palpable and has been expressed by several carriers previously reticent to build out these AI practices. Firms who have been slower to get their AI use case pipelines established expressed that they were intentional about not being "early adopters," choosing instead to be "fast followers." The challenge these firms have faced is that AI has demonstrated such explosive growth that "fast followers" — traditionally those who are only slightly behind any early adopters of technology — can risk becoming laggards. In terms of AI, the distinction between "early adopters" and "laggards" seems to be measurable in months, not years. About a third of the companies who chose to "let others go first" decidedly did so because they did not want to invite the risk, with two thirds stating that they were unable to redirect their technology resources and budget from ongoing digital transformations.

Under direction from their Boards and advisors to proceed, firms — especially smaller ones that do not have access to similar levels of capital and resources as larger companies — are trying to understand the best way for them to move forward. It is important to note that early in 2023, even when the "early adopters" invested in building out their AI use-case pipelines, most of the experiments were based on maturing existing AI implementations as well as employing a laboratory mindset for GenAI use cases. For the latter, these early-stage iterations were rudimentary and seemingly selected at random because they lacked a cogent narrative and strategy connecting them to delivering business value. A majority of these early experiments did not yield expected outcomes but gave the "early adopters" much needed time to build out repeatable processes and marshal the enterprise into developing the AI/GenAI use-case pipelines that they are benefiting from today.

A company does not always need to be an "early adopter" or "fast follower," but in the age of AI, being a laggard might mean falling well behind the industry. Carriers that have demonstrated their ability to be nimble and agile — and adapted to the shifting technology landscape — are likely to increase their odds of success.

If your organization is still in the early stages of considering or building out your GenAI use-case pipeline, the good news is that the costs and barriers to entry are generally minimal, and the ability rapidly scale up is easier than most other technology implementations. Note that this is true of GenAI, and not necessarily for other, more complex AI implementations such as those used for automated and accelerated underwriting, fraud detection, etc. Carriers generally have five options to jumpstart their GenAI journeys.

IN ORDER OF EASE OF ENTRY:

1. Carriers can leverage publicly available GenAI engines such as ChatGPT and DALL-E. This is the easiest way to explore the tools' capabilities. Note that the risk with public implementations always remains the loss of intellectual property.

2. A second option is for firms to partner with an AI technology provider. These include SaaS vendors and InsurTech firms. Note that the challenge here might be the time for the partnership to be established (contracts, etc.), and a broader challenge might be to discern how to integrate the vendor's offering into a carrier's existing technology ecosystem.

3. A third option that carriers can explore is to implement private versions of publicly available GenAI platforms.

4. A derivation of the third option is for carriers to establish private versions of publicly available GenAI platforms, and then customize them by allowing these private versions to be trained on a carriers dataset (along with the publicly available data that they come pretrained on).

5. Carriers can choose to build their own models from scratch. While this makes sense for complex AI models such as for underwriting, GenAI is quickly becoming a commodity and building your own large language models (LLMs) from scratch is likely unnecessary.

Maturing AI Use Case Centers: Centralization of technology efforts within a firm across our industry have been elusive due to the silos carriers tend to operate within. Centralization will be explored in more detail further in this study. AI, as an enabling technology, has likely been the first advancement toward which carriers have taken a decidedly centralization/center of excellence approach. This will be crucial to ensure that carriers are able to derive maximum benefit from their use case experimentation pipelines. It will also be important for leaders to ensure that they do not overswing the pendulum such that centralization erects artificial bureaucratic barriers that stifle innovation at a grassroots level.

Carriers in varying degrees of institutionalizing their AI centers (also called AI Governance Groups, AI Factories, Centers of Excellence, etc.) are generally following similar processes. An important first step has been to invest in educating a company's executives and employees on the safe and effective use of AI. Firms have developed their internal AI governance policies concurrently with this educational effort. More than two thirds of firms also blocked employee access to public GenAI systems such as ChatGPT and DALL-E at the same time to ensure that no sensitive information or intellectual property is accidentally placed in the public domain. Most carriers who have blocked access to public GenAI are working on implementing — or have implemented — private versions of these GenAI systems accessible only to the firm's employees. With a basic level of education established, carriers then built out their central AI teams comprised of individuals across the company's value chain. Overseen by the CIO, firms have then established an intake mechanism to vet and prioritize AI use cases. CIOs have been responsible for ensuring appropriate technology and infrastructure are available to support these prioritized use cases. It is expected that these AI centers will continue to mature as they measure the value being realized from these use cases and develop a pathway for operationalization.

Most carriers who have blocked access to public GenAI are working on implementing — or have implemented — private versions of these GenAI systems accessible only to the firm's employees. With a basic level of education established, carriers then built out their central AI teams comprised of individuals across the company's value chain.

One thing that leaders should be diligent about as use cases are being vetted is whether it is anticipated that a current technology vendor would deliver the same features from a particular use case under consideration. It will not be worthwhile to build a bespoke use case from scratch if a vendor would imminently offer the same features. For instance, carriers developed their internal document summarization engines. These GenAI systems can ingest vast amounts of data (including documents) and produce insightful summaries. However, these features are now an integral part of Microsoft Copilot, albeit at a cost. The cost of building and maintaining a bespoke set of features might not be worth it as the per-employee Copilot cost might be more cost-effective over the long term.

In addition to the use of AI in automated and accelerated underwriting, GenAI will deliver operational efficiencies, reduce costs, and increased productivity by being able to assist with the assessment of applicant risk and setting fair pricing by examining a vast range of disparate data sources.

Use Case Prioritization Maturity: Carriers will continue developing and refining how use cases are vetted and prioritized. This is a departure from what seemed to be a combination of existing use cases brought together from across a company with use cases that were being ideated upon at a "local" level. As GenAI continues to mature, so will the process by which carriers prioritize which use cases that leverage AI/GenAI should be prioritized over others.

Carriers will realize productivity enhancements, cost savings, operational efficiencies, automation, and increased means of customer engagement as the primary returns on GenAI investments. It is expected that carriers will continue to prioritize GenAI use cases that offer quick results. These include the build-out of private versions of public GPT, and leveraging turnkey GenAI that is now available in existing products such as Microsoft Copilot, Salesforce Einstein, Adobe AI, etc. Since GenAI can process and analyze a large corpus of data, it is likely to be beneficial in the underwriting process. In addition to the use of AI in automated and accelerated underwriting, GenAI will deliver operational efficiencies, reduce costs, and increased productivity by being able to assist with the assessment of applicant risk and setting fair pricing by examining a vast range of disparate data sources.

There has been significant focus on the ability of GenAI to deliver automation, cost savings, and operational efficiencies. However, GenAI can also drive growth and profitability by creating new products, augmenting existing product lines, expanding into new or untapped markets, unlocking new revenue streams, or helping expand existing ones. It is unsurprising that several GenAI use cases within companies today are applicable within IT departments. IT-developed GenAI tools help improve quality, accuracy, security, and scalability of code. Code generation AI tools are greatly expediting the software development process and improving developer productivity. A large number of GenAI use case are focused on realizing potentially significant cost savings and operational efficiencies. These are mostly low-risk use cases to generate new content in areas such as marketing for sales and branding, job postings and job description creations in Human Resources, requests for proposals (RFPs) in contracting, and a variety of low-risk, repeatable templates within corporate legal. Human judgment will still needed for these use case implementations.

GenAI will add value to Marketing and Sales/Distribution in customer discovery, acquisition, and engagement. By analyzing website data, and third-party sources such as social media data, GenAI can help create personalized product or coverage recommendations as prospective customers study insurance products. Most CRM platforms such as Salesforce and HubSpot will continue incorporating GenAI into their products. This, in combination with GenAI use cases within carriers, will help customer engagement by equipping agents and advisors with customer-specific information to allow for better engagement. This includes the ability to generate product and coverage summaries and real-time scenario planning and other illustrative tools. Similarly, in the worksite space, carriers will likely leverage GenAI to create customized insights and recommendations to employers and employees. By analyzing demographic information, health profiles, and historical claims data, GenAI can help recommend customized group insurance plans.

ACTIONABLE INSIGHTS

Balance the Need for Speed With Strategic Intentionality: Carriers have to have an AI strategy in place. This strategy needs to fixate less on the technology and more on business priorities and business problems that can be resolved by using AI and GenAI. To realize business value, AI and GenAI implementations cannot emerge organically. They need to be tied to direct business benefit and continued coordination across cross-functional teams. To truly get value out of GenAI implementations, firms are advised to provide GenAI training to employees. Reskilling employees is the key to lasting enterprise success with using GenAI to transform existing business processes.

While a significant number of carriers have made rapid progress on ideation with AI and GenAI, carriers are advised to continually ensure that the AI use cases are tied to tangible business outcomes, prioritized only after ensuring that common vendors such as Microsoft/Salesforce/Adobe/etc. are not planning to offer the same functionality in their products, or that turnkey alternatives do not exist in the marketplace (i.e., AI solutions being supplied by InsurTech-type firms). They should also "look left," that is, understand what adjacencies within our industry are doing with AI. It will be valuable for multiline carriers to evaluate AI developments within property and casualty. Retail banking can also be a good source of information on AI's directional progress, and avoid the urge to develop custom best practices, frameworks, and tools that can be best built by the industry for the mutual benefit of the entire industry. While carriers might achieve speed by operating alone, it will be vital to contribute to and learn from the LIMRA and LOMA AI Governance Group (AIGG). This will ensure sustained long-term success with AI programs for each firm.

To truly get value out of GenAI implementations, firms are advised to provide GenAI training to employees. Reskilling employees is the key to lasting enterprise success to transform existing business processes.

Manage Risks and Expectations: While AI/GenAI will transform the insurance value chain, it is important to ensure that carriers manage expectations — not all AI use cases will find success and it is likely that more than half of these use cases fail to achieve intended outcomes. That should be expected. However, firms would do well to focus on operationalizing the ones that do find success, since even if 1 in 10 use cases succeeds, it has the potential to create real and measurable business value. Firms should also carefully manage risks associated with GenAI implementations by establishing the appropriate guardrails. This can be accomplished by working with the LIMRA and LOMA AI Governance Group (AIGG) and customizing the governance-related outputs of this group for your specific company's needs.

There are many risks associated with GenAI and AI in general. The most prominent of these risks are those that deal with explainability and transparency. AI systems can be "black boxes" and the inability of an operator to clearly explain how AI arrived at a decision, can erode trust in the AI's output. Broadly speaking, the three primary concerns with GenAI are the loss of corporate intellectual property, cybersecurity and data/information security exposures, and reputational damage due to unexpected outcomes or bias. A prominent concern among all fields of AI is ensuring that systems are free of bias and do not cause proxy discrimination. GenAI models can tend to "hallucinate," which means that these systems can generate incorrect or misleading results. Hallucinations can be a result of one or several factors in unison. These include poor data used to train the model, an inadequate amount of appropriate training data, biases inherent in the data used to train the model, inadvertent correlations wherein the models make correlations that are due to unintended causation, incorrect model assumptions, etc. GenAI and AI systems present a new cyber risk threat landscape. By democratizing a bad actor's access to the same tools, GenAI can be used for creating sophisticated phishing attacks, deep fakes, etc.

State-level regulators are increasingly seeking oversight on insurers' AI models. The fluidity of regulation and regulatory frameworks is explored further on in the study. To mitigate these challenges, insurance companies should prioritize ethical issues over methodological issues when building AI solutions, establish robust auditing processes, implement governance models including those that will result from the LIMRA and LOMA AI Governance Group, stay abreast of ongoing regulatory developments, and establish means to derisk the vendor AI supply chain.

Broadly speaking, the three primary concerns with GenAI are the loss of corporate intellectual property, cybersecurity and data/information security exposures, and reputational damage due to unexpected outcomes or bias.

Humans AND Humanity in the Center: "Humans in the center" is a familiar refrain across AI to ensure that any AI decisions and recommendations are vetted by a person so that human judgment is always at the core. Most carriers are laser focused on this approach. However, in addition to this mindset, carriers should also ensure we maintain "humanity in the center" of AI implementations. Ours is a human-centered industry. At its core, it's about people helping other people protect their financial futures. While it is appropriate to invest in the use of chatbots and other AI for customer engagement and servicing, firms should ensure that they are striking an effective balance between employing AI and maintaining the human-to-human connection. Establishing empathy and exhibiting soft skills are crucial for more sensitive aspects of customer engagement and service, such as during the claims process.

Measuring Value: A Gartner report published in March 2024 comparing AI software investment across all lines of business in the insurance industry from 2023 to 2027 projects that "global AI software spending in the insurance market is forecast to increase 17.4% in 2024 to $9.5 billion and reach $15.9 billion by 2027, with a five-year CAGR of 18.2%."[1]

The ability to measure success of all AI implementations within a carrier in addition to the AI use case pipeline — at least in the near term — is going to pose a challenge for carriers. AI specifically has not been tracked as an investment item within IT budgets. Having business units budget for technology vendors that provide AI-enabled services (for example in underwriting or marketing) can exacerbate the challenge of having an enterprise-level view into AI spend. With technology vendors incorporating AI into their products, it will become increasingly challenging to isolate spending specific to AI within these product sets.

Measuring value and conducting cost-benefit analyses will be easier when it comes to the GenAI ideation and experimentation use case pipelines. If your use case pipelines originated organically — that is, they emerged within a business unit over the last 12 to 15 months and expanded across the enterprise — it is likely that these pipelines lack operational maturity and rigor, which means leaders need to be intentional about retroactively building a business process around these pipelines. The process should include not only what use cases get prioritized, how they get vetted, and who would be responsible to shepherd them, but they should also account for what soft and hard costs associated with these experiments might be, and more importantly, what the efforts, costs, and expected measurable outcomes associated with operationalizing these use cases might be. Leaders should be intentional in defining objectives that are associated with desired outcomes for each use case. These objectives should align with the broader enterprise priorities and include basic objectives that one would expect from any technology initiative. For example, reduction of costs, operational efficiencies measured in FTE, productivity lifts, accuracy, and quality of IT code, etc. Leaders should consider building these KPIs and represent them on management dashboards/scorecards such that they can be socialized with the senior leadership team and Boards. The LIMRA and LOMA AI Governance Group will be developing turnkey CBA/ROI tools that carriers can employ (or customize) for their needs.

[1] Compare AI Software Spending in the Insurance Industry, 2023-2027, Gartner Research, 2024.

RAPID READ:

The establishment of enterprise AI governance is vital for the success of AI and GenAI implementations. The AIGG explored regulatory frameworks currently being employed within carriers, firms' internal preparation for potential regulation, vendor supply chain risk management, mitigation strategies for bias and discrimination in AI implementations, and the need for AI explainability and transparency. Firms are advised to adopt a human-centric approach to AI strategies. Organizations that fail to build a robust governance model will face challenges in scaling their AI programs in the future. A shared enterprise definition of AI risk and the incorporation of risk management into AI strategies will be vital in ensuring AI success over the next two years and beyond.

3.2 Governance

The implementation of AI within organizations warrants more attention to all aspects of governance than any other technology. The topic of AI governance is multifaceted, covering everything from how organizations are planning to keep abreast of regulations and/or regulatory frameworks to how carriers are structurally preparing internally for effective execution of established governance policies, and from how to derisk the vendor supply chain to how carriers are managing the sensitive topic of ensuring that their AI implementations are free of bias and proxy discrimination, and how to adhere to the tenets of AI explainability and transparency. Absent industry-specific regulation, carriers are preparing for potential regulation or regulatory guidelines at the state level. AI is evolving at a pace that regulation has not been able to keep up with, and it is unlikely the United States will have a Federal AI regulation (akin to the European Union (EU) and the EU AI Act) in the near future.

Overall, responses from the AIGG highlight a strategic but varied approach to AI governance, with companies at different stages of implementing policies, risk management strategies, employee education, and ethical considerations. The emphasis on structured governance models, the proactive establishment of risk profiles, and the initiation of educational programs are reflective of an evolving landscape of AI governance. However, the cautious approach to transparency, explainability, and bias indicates ongoing challenges in aligning AI initiatives with ethical and regulatory standards, especially where those standards are yet to be fully defined.

A majority of firms with ongoing use case pilots have appropriately established governance models that incorporate AI-specific policies and procedures, including aspects like ethical use, data privacy, and cybersecurity. This study loosely structured governance models as the "4Ps" — policies, protocols, procedures, and processes. Firms are actively managing all aspects of AI governance through their AI task forces/

governance groups. These centralized cross-functional groups are overseeing all aspects of governance over enterprise AI initiatives. There's a mixed approach to building AI risk profiles (cybersecurity, data privacy and protection, regulatory and compliance), with some companies actively monitoring and preparing for AI regulation through legal and compliance teams, while others are still in the early stages of this process. This indicates varying levels of readiness for future AI regulation and associated risks.

The state of employee education on AI varies. Some companies in the early stages of creating AI literacy programs, while others have not yet begun their education efforts. This suggests that while there's recognition of the need for AI literacy, the implementation is uneven. Companies are aware of and emphasize the importance of strong vendor management programs to mitigate risks associated with AI. This includes the evaluation of AI technologies and their providers, though some express concerns about the adequacy of current practices. Responses indicate a concern about vendor management, especially regarding the control and transparency of AI solutions provided by third parties. Some companies have initiated vendor management programs to address these concerns, reflecting the complexity of integrating external AI solutions.

Companies are taking steps to ensure AI transparency and explainability, with some establishing task forces focused on these issues. However, approaches vary, with some companies prioritizing regulation compliance before expanding efforts, indicating a cautious approach to explainability. Most firms expressed a recognition that more needs to be done, especially in regulated areas or where explainability is critical. Preventing AI bias and discrimination is a recognized challenge, but many companies admit to not being there yet or avoiding risky areas where explainability and bias could be significant issues. Some are focusing on building in-house AI expertise to tackle these challenges, while others are intentionally avoiding use cases that could raise discrimination concerns. Responses indicate a varied approach to mitigating bias and discrimination in AI, from acknowledging the challenge and taking initial steps to avoid risky areas, to not yet undertaking significant actions in this direction.

The state of employee education on AI varies. Some companies are in the early stages of creating AI literacy programs, while others have not yet begun their education efforts. This suggests that while there's recognition of the need for AI literacy, the implementation is uneven.

IMPLICATIONS

Build a Holistic, Trusted AI Governance Framework: Carriers have done well in building internal governance policies associated with their GenAI use case pipelines. Overseen by the internal governance groups/task forces, carriers recognize the need for a risk-based approach to implementing and scaling AI solutions that traverses the enterprise. It will be imperative to develop a holistic AI governance framework that covers all ongoing AI activity within a firm, as carriers coalesce their disparate ongoing AI activities under a centralized body. Management of AI risk should be across the enterprise value chain and carriers are advised to be diligent to ensure they mitigate inadvertent risk transference within a company, wherein implementation of a perceived governed AI solution does not create or expose risk in a different part of the value chain.

For most organizations, the establishment of governance around GenAI use cases is a good starting point, but carriers will need to centralize it for any pre-existing AI-enabled programs, from underwriting to marketing, etc. In the absence of an industry governance model, the LIMRA and LOMA AI Governance Group will be publishing best practices and guides for the entire industry to follow. These can be extended and customized to an individual carrier's needs and will cover AI holistically (not just focusing on GenAI). When these best practices are released to the industry, companies would do well to integrate/apply them to any governance frameworks they have developed. In the interim, carriers should ensure that their governance models encompass a broad variety of risks, and are predicated on factors that include prioritizing AI explainability and transparency, basing AI programs on ethics, the mitigation of bias and proxy discrimination, data and information security, data privacy and protection, cybersecurity, promoting continual oversight and monitoring, and assignment of responsibilities (RACI), etc. Carriers should ensure that ethics remain at the core of AI development, use appropriate data for training (fit for purpose), promote transparency and explainability, and implement robust audit, quality control, and quality assurance throughout the development process.

For most organizations, the establishment of governance around GenAI use cases is a good starting point, but carriers will need to centralize it for any pre-existing AI-enabled programs, from underwriting to marketing, etc.

Organizations that fail to build a robust governance model in 2024/2025 will be highly challenged to scale their AI program and GenAI use cases in 2025 and beyond. However, it is also true that organizations overengineering their governance frameworks and ceding enterprise agility in bringing use cases to operationalization risk falling behind very quickly. Leaders have to adopt a risk-based approach that balances the need to avoid enveloping AI use cases in unnecessary bureaucracy. Leaders should expect this balance to be a challenge and

require constant diligence to ensure that the scales do not tip in the direction of agility without guardrails or conversely, in a direction that stymies innovation. Carriers should keep the "humans in the center" approach in mind. Human domain expertise is absolutely necessary for development, adoption, scaling, testing, auditing, and attestation.

Manage Existing AI Risks: As companies develop their governance frameworks, they should take a risk-based approach. Primary among these risks is that AI systems can often be considered "black boxes," preventing us from understanding how a system rendered a particular decision. Trusting the output from these opaque AI systems is impossible without an understanding of how the system arrived at the output or recommendation that it did. The challenge is exacerbated by the fact that AI systems can make potentially trillions of computations, making it impossible for humans to follow the decision pathways and begin to comprehend the output. AI models can also inadvertently establish incorrect relationships and causality. Carriers should focus on ensuring Explainable AI is at the core of their governance models because the results of making business decisions based on incorrect or flawed outputs could be devastating. Addressing concerns with bias and proxy discrimination should be paramount.

While Managing Additional GenAI Risks: One of the greatest concerns with GenAI has been the loss of corporate intellectual property. The ability for IP and copyrighted materials to show up in GPT engines remains high and since GenAI has no citations or attributions, it opens a challenging set of circumstances for IP infringement claims. Data privacy and protection continue to be ongoing concerns within all facets of AI, but if firms use public GenAI systems, the loss of identifiable customer information into these engines could result in severe data privacy issues.

ACTIONABLE INSIGHTS

A Shared Enterprise Definition of AI Risk: The value proposition that AI centralization has unlocked within enterprises has also presented firms with an opportunity to manage the totality of risk related to AI implementations across the company. Firms will have to report to their internal and external stakeholders on AI risk holistically and the only way to do that is to manage it at an enterprise level. Until recently, AI risk management had been relegated to operating within the context of ongoing AI programs. These disparate implementations stretched Chief Risk Officers, Chief Information Officers, Chief Information Security Officers, etc. with having to individually assess and manage risk for each application of AI. Cross-functional AI teams — task forces, governance groups, centers of excellence, etc. — will be critical in driving AI risk management at the enterprise level. These groups will have a 360-degree perspective of the totality of enterprise AI risk. This is something that has eluded companies, since it was implausible to expect any single department to have had a holistic understanding of the risk, how it connects across the value chain, and effective risk management strategies. Defining roles and responsibilities across these horizontal cross-functional groups, and how these groups liaise vertically across individual departments

will be of paramount importance. Leaders should guide these groups to perform another key function that has evaded firms thus far — the need to establish a common taxonomy across the enterprise on what is considered a risk within the context of AI. With most major technology providers incorporating GenAI capabilities into their products, AI will be ubiquitous within operating systems, browsers, Microsoft, and Google productivity software (from Microsoft Word to Google Docs), etc. Defining and then articulating risk within these products will prove challenging. It will be important that the entire company have a common way of identifying these risks before being able to manage them.

Incorporate Risk Management Into AI Strategies: This study has made a strong recommendation that whether a firm builds its own AI capabilities or purchases them from a third-party technology vendor, it is vital to have an AI strategy. Whether you leverage AI for complex functions such as underwriting or are intending to leverage GenAI, your organization needs a strategy. Even if you decide to not pursue AI/GenAI, GenAI will be in your firm by virtue of the everyday technology products your employees use today. AI risk management should be a central part of this AI strategy.

Companies should outline a two-year roadmap of their products, services, and business processes that will likely leverage AI/GenAI, either to augment or reengineer/reimagine them. These AI strategies should outline risks associated with each type of AI implementations (for instance, AI for underwriting, GenAI for code authoring, GenAI for marketing materials, etc.). These risks would likely fall into the familiar categories of data and information security, data privacy, AI model risk management, compliance, third-party vendor management and derisking the AI supply chain, and operational risks. Once an inventory of risks has been outlined, carriers should rank these risks. Typically, ranking AI risk would depend on several factors, including but not limited to customer-facing systems, reputational damage risk, financial risk, violations to code of conduct, regulatory and compliance issues, customer impacts, novel data sources, use case complexity, use in developing new products, opacity of third-party vendor solutions, etc. Once risks have been identified, defined, and ranked, carriers can then develop controls/compensating controls against each risk. Low-risk AI implementations will not require the same amount of scrutiny, oversight, and controls as higher-risk implementations.

Implement AIGG Outputs: The LIMRA and LOMA AI Governance Group will be developing industry best practices around multiple facets of AI governance. This includes but is not limited to an acceptable use of AI across the insurance value chain that is stylized to the EU AI Act. For multinational carriers, compliance with the EU AI Act will be mandatory. The AIGG will also prescribe AI governance models, as well as AI maturity models. Note that these best practices will be generic to all carriers across the industry. It is likely that smaller carriers will be able to adopt these best practices in a turnkey manner. All carriers are free to extend and customize these frameworks to suit their specific needs. Ideally, carriers who already have developed their own governance models will be able to co-op best practices from the output of the AIGG, and combine them with their own models. The AIGG outputs will prescribe, but not delve into, the state of data management within each carrier. It will be vital for firms to implement robust data governance, especially for data that is used for AI purposes. Data management will require carriers to understand all aspects of data being used for AI: the origin of the data, lineage, any transformations that occur along the way, quality, completeness, etc.

Companies should outline a two-year roadmap of their products, services, and business processes that will likely leverage AI/GenAI, either to augment or reengineer/reimagine these products and processes. These AI strategies should outline risks associated with each type of AI implementations (for instance, AI for underwriting, GenAI for code authoring, GenAI for marketing materials, etc.).

LEARN MORE about the AI Governance Group (AIGG) **here**.

Advancing the financial services industry by empowering our members with

KNOWLEDGE

INSIGHTS

CONNECTIONS

SOLUTIONS

LIMRA LOMA

Navigate With Confidence

0699-2024