



Deploying Agentic AI Across Insurance and Retirement Products

Anumandla Mukesh

Independent Researcher

mukeshanumand@gmail.com

Abstract

The integration of agentic AI into enterprise data products, following established principles of strategic decision-making, significantly improves agent-in-the-loop decision-making within insurance and retirement sectors. Decision-making systems underpin a broad spectrum of enterprise intelligence solutions, enabling human decision-makers to harness advanced AI technology, data analytics, and internal expertise. However, agent-in-the-loop solutions must address the same set of decision-making challenges that AI systems tackle independently—commanding consideration during the architecture design phase, especially in data products integrated for concurrent consumption within enterprises.

Within enterprise data products, agentic AI in support of strategic decision-making exposes tertiary product dimensions that determine the product's trustworthiness, its behavioral alignment with user expectations, the underlying technical infrastructure, and the managed change required for adoption and operation. It serves as a lens for examining enterprise data products structured for concurrent syndication, spotlighting product characteristics that ensure decision-making processes yield agent-in-the-loop systems of record supported by AI-assistive technologies and the agents' governing processes. The approach follows convergence techniques mirrored in industries such as cybersecurity that also underpin enterprise data products serving strategic decision-making in the insurance and retirement spaces.

Keywords : Artificial intelligence · Strategic decision-making · Insurance and retirement · Enterprise data product · Data infrastructure · Agentic AI.

1. Introduction

The analysis situates enterprise data products at the intersection of agentic AI and strategic decision-making in the insurance and retirement solutions domains. Artificial intelligence (AI) capable of operating independently in the world is increasingly seen as a cornerstone of data-driven automation and real-time, contextually aware decision-support systems. Such agentic AI adds human-like cognitive capabilities to enterprise data products, amplifying their impact by facilitating human-machine collaboration. Despite the associated emergence of enterprise data products for supervisory decision-making—from lowered risk alerting in risk assessment, underwriting, investment management and pricing to heightened personalization in product recommendations and pricing—these products have yet to be holistically conceptualized in these solutions spaces. Therefore, the study integrates strategic human decision-making processes and enterprise data products into a single framework, explicitly delineating their various components and interactions.

Many organizations subject to extensive regulatory supervision collect vast amounts of data to respond to alerts and recommendations generated by enterprise data products—often labelled under supervisory decision support or explainable AI—and yet these alerts are rarely heeded. Maturity models focused on real-time, agentic data-products in these solution spaces remain limited. The integration of agentic AI and enterprise data products could represent a new frontier of impact created



through an organization's data assets. Enterprise data-products augment human decision-making by providing supervisors contextually relevant information at appropriate times and can govern the actions of automated decision and process systems. The ability of such systems to simulate and predict—along with a robust and trusted architecture—opens the door to deployment in more ‘human’ processes, where supervisor engagement (and faith in the systems offering insights) can be fostered. Data-products that enhance strategic supervisory decisions thereby allow a leap in operational trust that is a key missing piece to realizing significant operational automation.

1.1. Background and Significance

No single product, not even the hallowed internet, will undergo as profound a reconfiguration this century as the enterprise data product. A more general, widely applicable concept than a specific product category, enterprise data products are the conduits through which organizations mobilize their data assets to create sustained, quantifiable, data-driven decision-making improvements. These products will continue to feature data-aware rules and code checks, analytics, and business intelligence dashboards. Over time, however, enterprise data products will increasingly incorporate synthetic data, domain-specific motion planning, and actuation capabilities. Enterprise data products are naturally poised for this evolution because of the central role they play in delivering specialized and bespoke decision-making improvements.

Agentic artificial intelligence (AI)—defined as systems that act for themselves, in an environment that combines play and eeriness, and with the attitude of a human formal agent—is becoming a hot topic. Agentic LLMs powered by Internet AI access current data and sufficient computer time have displayed some emergent abilities typically associated with higher mental aptitude. Their functionality position them somewhere between a chatbot and a software interface. A narrow class of enterprise data products has appeared in this space, but it is still an early-stage development. The addition of an agentic agent to an enterprise data product adds a fantastical dimension and promises further specialized and bespoke decision-making improvements.



Fig 1: Agentic AI for Enterprises

1.2. Research design

Research in the insurance and retirement domains focuses on enterprise data products that address strategic business decisions. Employee resourcing, auditing, and governance are managed without the risk posed by agentic AI systems. These products



almost exclusively consume the data. Integrating agentic AI systems delivers full functionality across core applications and supports product innovations, such as education, gaming, or analytics.

The proposed framework incorporates the full range of enterprise data products in a risk-mitigated manner that extends trust, safety, and governance mechanisms from human users to machine agents. Elements supporting these features rely on external data management, security, or compliance tools that are well understood and mature in enterprise deployments. Use case scenarios cover foundational business functions—underwriting and risk assessment, and product personalization and pricing—whose common motivations are greater responsiveness to individual customer needs and larger revenue streams.

2. Conceptual Foundations

Enterprise data products that incorporate agentic AI can facilitate diverse forms of strategic decision-making in the insurance and retirement sectors. These segments involve assessment of forward-looking events and complex structures, introducing information asymmetries and uncertainties that determine insurability and value-at-risk. Agentic capabilities enable users to explore future scenarios or system changes that may alter positions for better or worse, allowing Windows of Growth, Harvest or Professional Survival to be identified. Decision-makers targeting these outcomes can seek relevant data-driven insights.

Complex decision-making processes are involved in these markets, linking the integration of agentic AI into enterprise data products to well-defined decision objectives of C-suite executives, business owners and product managers. Decision objectives for CEOs and CFOs of insurance and pension companies include achieving a sustainable return-on-capital-outlay through optimising investment of equity and reserves, managing longevity exposure within risk appetite while maintaining required reserves and enhancing market penetration without compromising strict capital requirements. For business owners of pensions, the focus is on the trade-off between growth and risks arising from investment strategies. For product managers, the essential task is to identify changes in market needs or regulatory framework that require the design or modification of products (for pricing, coverage, etc., making them more relevant to target markets). Strategy frameworks that facilitate improved decision-making generally adopt a time outlook of one year or more.

Equation 1: Derived equations (step-by-step, from the paper's constructs)

- Enterprise data product generates “a set of inputs” for the agentic AI.

A Framework for Integrating Age...

- Inputs are assembled from enterprise-ready sources plus additional sources (source/support data).

A Framework for Integrating Age...

Step-by-step derivation

1. Let raw data sources be split into:
 - D_s = source data layer (non-product sources required)



- D_{sup} = support data layer (operation-supporting sources)
- 2. Let the enterprise data product (EDP) apply:
 - selection, cleaning, joining, standardization, lineage tagging (the paper emphasizes provenance, quality, catalogs)

A Framework for Integrating Age...

- 3. Model that as a transformation function:
- 4. $X = f_{EDP}(D_s, D_{sup})$
- 5. where X is the enterprise-ready dataset used as the agentic model input.
- 6. Agentic AI consumes X and produces interpretable outputs in the user experience layer:

$$\hat{y} = f_{AI}(X)$$

2.1. Agentic AI and Enterprise Data Products

Agentic artificial intelligence (AI) is capable of concurrently joining decision-making, execution, and active learning within the context of enterprise data products. Supported by evolving data, AI, and software engineering capabilities—including secure, privacy-preserving data access; high-quality data with known provenance; specified permissioning policies; and natural language interfaces (NLIs)—agentic AI expects to become the de facto architecture for strategic data science applications.

Established production-grade enterprise data products—bespoke data assets that deliver optimal business impact and are used by multiple teams across the business—typically lack traffic policies tailored to the user journey of risk and life underwriting. They partition analytics responsibilities across the complexity spectrum without blending decision-making, execution, and dynamic machine learning coupled with production-grade automation. Consumer data products have also to date only patchily explored connecting corporate or customer actions with decisions or recommendations generated by the data product.


2.2. Strategic Decision-Making in Insurance and Retirement Domains

Strategic decision-making in the domains of insurance solutions and retirement solutions is characterized by high complexity, infrequent repetition and a direct impact on long-term profitability and shareholder value. It benefits from decision-enhancing capabilities rather than full automation. Insurance and retirement solutions experience cyclical contention and conflict, between their roles as value protectors to customers during their time of need and as profit-generating streams in a capital-intensive industry.

Breach events – be they underwriting decisions at the time of an insurance policy application, the terms of a financial product, or regulatory approval of insurance companies for systemically relevant risk – generally get blamed on the industry for technological deficiencies or immoral market pricing strategy. Yet biases in data, explainable reasons behind the decisions taken, fact-finding and mitigation of the root causes leading to the breach event tend to fall under the conditions of explainable artificial



intelligence, requiring intent alignment and decision provenance for dependability. Stakeholders, however, are concerned about the determinism, perceived excessive basing on logical-dependency-rule chains for detection with false positives, false negatives and brittleness with respect to novel or rare instances.



3. Framework for Integration

Enterprise Data Products are custom-built technology platforms that the business owns and uses for executing its operating strategy and for making decisions that shape and guide the organization's business strategy and/or operational strategy. They are built on consistent and reliable provenances for selected data from across the enterprise and allow for enhanced query



performance, in-memory processing, heuristic-based algorithms, greater processing capacity for large volume data sources, and/or the availability of domain expertise and knowledge to consumers at speed and scale. Such products can be developed for a specific decision, or for a broader set of decisions with a more encompassing data environment. Enterprise Data Products remain the central component of the enterprise data ecosystem, together with the data that organizations are willing to store in a shape, structure, design and content format that support decision alignment.

There are specific types of Enterprise Data Products for Insurance and Retirement. For an Insurance Enterprise Data Product, its architecture is designed to make changes happen more quickly and for the organization to monitor operating performance in real time. It may include an Enterprise Data Product for Underwriting, an Executive Dashboard that presents Insurance Product Data and MI, a Business Operations Monitoring MI Dashboard for Operation Services, and an Insurance Data Mart that enables multidimensional analysis. For a Retirement Enterprise Data Product, put simply, an Enterprise Data Product for Retirement is what the Retirement business uses to execute its operating strategy in shaping and responding to market demand.

3.1. Architectural Overview

For decision-making processes requiring the consideration and analysis of numerous factors, an enterprise data product should be conceived as a generator of a set of inputs to be provided to an agentic AI model. Miller and colleagues propose a framework for developing an enterprise data product with agentic AI. It consists of an architectural overview depicting how an agentic AI can be integrated into the enterprise data product architecture, the data infrastructure and provenance required, illustrative use case scenarios in the insurance and retirement domains, and relevant trust, safety, and governance mechanisms. An agentic AI product is an agentic AI model with access to an output of an enterprise data product. Such a model operationalizes a role requiring analysis, reasoning, prediction, recommendation, evaluation, or similar processes in the insurance and retirement domains.

A conceptual architecture to support the integration contains five layers: user experience, enterprise data product, agentic AI, source data, and support data. The user experience layer consists of the user interfaces and interactions of the enterprise data product and the decision-maker-facing interpretable agentic AI outputs. The enterprise data product layer contains enterprise-ready data sources corresponding to the inputs required by the agentic AI. The agentic AI layer consists of the model to be integrated with the enterprise data product. The source data layer encompasses any data sources that provide the source data required by the enterprise data product or agentic AI model but are not an output of the enterprise data product. The support data layer constitutes data sources that assist the proper functioning of the enterprise data product or agentic AI model.

For complex decision-making processes that involve weighing and analyzing many interdependent factors, an enterprise data product can be conceived as a structured generator of decision-ready inputs for an agentic AI model. Building on this idea, Miller and colleagues propose a framework for developing enterprise data products specifically designed to integrate with agentic AI, combining an architectural overview, data infrastructure and provenance requirements, illustrative insurance and retirement use cases, and the trust, safety, and governance mechanisms needed for responsible deployment. In this framework, an **agentic AI product** is defined as an agentic AI model that can access and operate over the outputs of an enterprise data product, effectively operationalizing domain roles that require analysis, reasoning, prediction, recommendation, evaluation, or similar decision-centric capabilities. The proposed conceptual architecture is organized into five layers: **(1) user experience**, which includes the decision-maker interfaces, interaction workflows, and interpretable AI-facing outputs; **(2) enterprise data product**, which packages enterprise-ready data aligned to the AI's required inputs; **(3) agentic AI**, representing the model and its embedded decision logic; **(4) source data**, covering external or upstream datasets required by the product or model but not produced as part



of the enterprise data product; and (5) **support data**, which supplies auxiliary resources such as reference datasets, metadata, or operational signals that enable reliable functioning of both the enterprise data product and the agentic AI model.




Fig 2: Architectural Overview

3.2. Data Infrastructure and Provenance

Enterprise data products with agentic AI provide a rich foundation for business users to discover and curate information that serves as a basis for richer predictive and analytical capabilities. Enabling these capabilities requires substantial investment in ensuring a robust, scalable, secure and agile data infrastructure that strives toward completeness, thoroughness, accuracy, quality, provenance and richness. These attributes are particularly important for data that is being used for regulation or compliance-related purposes, for which support for legal admissibility and auditability is likely to be necessary. Data must therefore be discoverable through high-quality data catalogs, support provenance tracing to establish the context in which the data was generated, and provide sufficient information for regulatory authorities to verify any internal verification—they should a fully-fledged air-gapped system, as there would always be a risk that such an internal audit could be faked by the business units. The technical implementation would also have to be supported by a dedicated governance model involving change management processes with involved business unit dealing with senior stakeholders and line of business heads.

Regulatory requirements demand a thorough underpinning for any data being used for regulatory-related purposes, including the marts feeding the enterprise data products. Dedicated data quality rules supported by tooling must be in place to establish trust. A full set of data quality metrics, measuring rules, tooling and governance setup would therefore be needed. If there are large and complex enterprise data products having a control over the internal layer already mature and the mapping into regulatory data



formats for reporting purposes is too manual, one option to consider would be to build a focused dedicated team that develops Maker-Checker-type models or similar.

4. Use Case Scenarios

Integrating agentic AI capabilities with enterprise data products creates an opportunity for augmenting decision-making in domains supporting long-term customer engagement, such as insurance and retirement solutions. Two application areas are examined in more detail. The first relates to underwriting and risk assessment within a societal context where increased climate risk is impacting populations around the globe, especially in regions unable to support adequate infrastructure development. The second focuses on product personalization and pricing to meet the growing expectations of some customer segments.

In the area of underwriting and risk assessment, there is a current tendency for insurers to silo climate-related data into separate datasets and models focusing specifically on climate risk. For example, Swiss Re's Net-Zero Approach uses a set of proprietary climate risk models to assess the impact of climate risk on the Solvency Capital Requirement (SCR) for each unit of the company and supports a Net-Zero underwriting standard by asking specific climate questions in reinsurance submissions. Such approaches, especially when coupled with the growth in catastrophic modelling, totally ignore the interplay that can occur between climate risk and other classes of risk undertaken by underwriting teams, something highlighted recently by Swiss Re Using the enterprise data product, climate risk data can be layered on top of the rest of the risk data for each underwriting submission and help decision-makers understand how climate risk contributes to the expected losses at a portfolio level. In addition, the potential of applying climate risk data to different asset classes (primary insurance, credit, operational) to inform underwriting considerations and support risk assessment can also be considered.

Equation 2: Governance constraint equation: “recommendations allowed only if auditability + regulatory alignment are satisfied”

Step-by-step derivation

1. Let:
 - $A(\hat{y}) \in \{0,1\}$: auditability test passes?
 - $R(\hat{y}) \in \{0,1\}$: regulatory alignment test passes?
 - $E(\hat{y}) \in \{0,1\}$: explainability threshold met? (explanations needed for trust)

A Framework for Integrating Age...

2. Define a “deployment gate”:

$$G(\hat{y}) = A(\hat{y}) \cdot R(\hat{y}) \cdot E(\hat{y})$$

3. Enforce action rule:



$$\text{Execute}(\hat{y}) = \begin{cases} 1, & \text{if } G(\hat{y}) = 1 \\ 0, & \text{otherwise} \end{cases}$$

4.1. Underwriting and Risk Assessment

To ensure the conduct of adequate underwriting and risk assessment, insurance companies must authenticating customers with high confidence and assess the risk they represent for the insurance company. An agentic AI can assist with these tasks by analysing behavioural events (i.e., transactions) and correlate them with pre-existing events in order to determine whether the event was fraudulent or trustworthy and, thereafter, predict if the event opens an opportunity for the insurance sector.

With respect to risk assessment, predictive analytics and agentic applications allow to analyse external data expecting to early validate a loan and change previous conclusions on the origin and possible destination of money concerning the study (i.e., to be discounted or not). Such approaches avoid entering the dirty money in the banking sector, a bottleneck in the financial system due to a disproportioned number of money that are washed.

4.2. Product Personalization and Pricing

Insurance and retirement solutions must cater to a diverse range of customers with distinctive preferences, behaviors, and risks. Enterprise data products can leverage past customer data to identify segments and personas, use agentic AI to actively learn from new customers, and match them to personalized, high-value offerings. Such dynamic and continuous product personalization requires incorporating deep generative multimodal models to create new promotional marketing assets with much higher conversion rates. Enterprise data products that combine past data analytics with agentic AI can be instrumental in deriving appropriate pricing for product segments and segments targeted by personalized marketing promotions to achieve profitable sales.

Pricing of insurance products involves significant work, usually requiring large teams, extensive market research, and annual cycles. With past insurance data in the data lake, enterprise data products can recommend suitable pricing for segments based on past segment profits, leveraging agentic AI to continually learn and recommend new price points. These pricing recommendations can be augmented by new pricing-generating capabilities based on active customer personas in Zone 2. In personal insurance, for example, a multimodal generative model can analyze home and home contents photographs uploaded to the social media by new customers in Zone 1 to create new marketing and activity-guiding assets for their next customers—step-by-step instructions for help in home protection or avoidance of colds or flus during the winter.

5. Trust, Safety, and Governance Mechanisms

Incorporating agentic capabilities into enterprise data products creates several safety, trust, and governance challenges. Explainability and transparency of AI model outputs are crucial to instilling business user trust in the decisions driven by those outputs. Business users need to verify the internal logic of underlying AI model parameters and choices, comprehending the development, training, and validation data quality and provenance. Audit capability is needed to demonstrate that risk and regulatory constraints are met. Such capabilities are especially important in the insurance and retirement solutions domain, where the consequences of business decisions made on AI model outputs could be catastrophic if implemented in an unsafe manner.



Meeting these challenges requires the development of appropriate support processes and controls, which may be encapsulated in an AI governance framework. Stakeholders throughout the enterprise can align on requisite trust, safety, and compliance characteristics of agentic AI products—both basic requirements and those emerged through the test-deploy-evaluate cycle. These characteristics thereby enable suitable oversight, reduce risk, and assure compliance with necessary laws and regulations from governing authorities such as the European Union and the United States.




Fig 3: Trust, Safety, and Governance Mechanisms

5.1. Explainability and Transparency

The implementation of agentic AI within enterprise data products requires mechanisms that ensure decision-making is explainable and transparent to users. These mechanisms facilitate downstream processes by providing individuals with the necessary contextual information and rationale to assess and act on potential outcomes. The preference among AI adopters worldwide for human-understandable explanations is strongly supported by literature showing that human decisions generally benefit from detailed decision models describing what constituents were used in determining the outcome or prediction. For underwriting, risk assessment, product personalization and pricing, the adequacy of product, pricing and personal situation for the policyholder can be understood by non-experts in insurance and thus become a key step in the explainability of agentic AI.

The aggregation of diverse internal and external data across multiple regions enables a comprehensive view of the customer's personal, actuarial and financial situation, necessary for developing a retirement, health, life insurance or general insurance strategy. Further personalisation can be accomplished by deep understanding of the accident frequency and severity likelihood by product type, which can be transformed into frequency-severity trend maps that show the possibility of accidents of all types, allowing for the definition of an adequate product offering. Transparency is different from explainability but is also necessary. As AI models become deeper, more complex and more multifaceted, they grow less accurate and require large datasets with different viewpoints. Agentic AI creates a multi-system view that needs to be shared among all agents. In this context, enterprise data products enable the inclusion of agentic AI and the consolidation of trust, safety and governance mechanisms for agentic AI use.

5.2. Auditability and Regulatory Alignment

The increasing use of agentic AI in enterprise data products is fraught with challenges regarding auditability and transparency. Addressing these challenges is essential for ensuring responsible, trustworthy, and safe use of generative AI systems in insurance and retirement solutions. Such systems must be capable of justifying recommendations and suggested actions related to the underwriting and pricing of insurance products, so as to provide sufficient insight into the internal decision-making process. This



is critically important not only for auditability, but also for regulatory alignment, particularly in the context of underwriting, risk assessment, and other business functions that require specific government approvals, external governance, or regulatory interventions. Action should be taken in line with prevailing regulations.

A supportive technical architecture must be established to facilitate examination of agentic AI decision-making processes, integrate agentic AI insights with knowledge graphs, and enhance the explainability and auditability of enterprise data products. Depending on the organization and the business function concerned, explainability and auditability of agentic AI systems must be taken on directly by enterprise data products, or enabled through integration with the underlying data infrastructure. Relevant policies and guidelines, clarifying the extent to which agentic AI-enabled voting decision-making processes should explain their suggestions, predicting key features that contribute most to voting decisions, and ensuring that rules-of-thumb based on training data remain in alignment with the steered knowledge sources, should also be established in conjunction with technical capabilities.

6. Implementation Roadmap

An integrated maturity model, consisting of successive phases, facilitates the implementation of the proposed architecture. Insurance and retirement solutions companies are traditionally based on actuarial principles. However, insurance solutions, which allow customers to address extreme but relatively rare consequences of various life scenarios, and retirement solutions, focused on developing long-term saving habits, are often subject to failures affecting a small proportion of customers, creating reputational risk. As a consequence, agents do not fully convince customers of the product's protective function. Agentic AI technology opens up the possibility of adding more advanced features to data products, with definite advantages in terms of satisfying customer needs in the purchasing process of appropriate insurance products such as life insurance, health insurance or pension savings.

The integration of agentic AI-enabled and enterprise data products requires advanced change management and stakeholder engagement processes across multiple phases. To facilitate customers' acceptance of underwriting decisions, a maturity model that encapsulates the main processes and technological aspects is required. A road map aligned with the initiative is also necessary.

6.1. Maturity Model and Phases

A maturity model provides a structured framework to assess the degree of organizational preparedness for implementing agentic AI capabilities and guides the gradual expansion of such capabilities within enterprise data products. The model consists of five maturity levels, each proposing a set of guidelines that assist enterprise data product stakeholders in systematically assessing and implementing the integration of agentic AI capabilities into data products.

At Level 1, enterprise data products are solely designed for anticipatory use cases, catering to passive stakeholders such as business analysts and middle management. Application and decision execution, including continuous monitoring, are solely handled by human users and therefore fall outside the product scope. The defining success criteria concentrate on the volume and quality of users' exploratory analyses. Level 2 enterprise data products are equipped to support agent-assisted scenarios that enhance business outcomes through user-AI collaboration but driven by human stakeholders. Such enhancements may include a



wide variety of features or suggestions, ranging from product-specific explorative queries to product-agnostic plans that restrict the search space to promising regions. Decision execution remains exclusively in human hands.

Equation 3: Data quality + lineage readiness equation (because “quality + availability + provenance” are fundamental)

A Framework for Integrating Age...

Step-by-step derivation

1. Define normalized scores in $[0, 1]$:
 - q = intrinsic quality (accuracy/consistency/timeliness)

A Framework for Integrating Age...

- v = availability/volume/on-target delivery

A Framework for Integrating Age...

- p = provenance/lineage traceability (catalog + provenance tracing)

A Framework for Integrating Age...

2. “Readiness” should fail if any critical component is near zero, so use a multiplicative form:

$$\text{Readiness} = q \cdot v \cdot p$$

3. Introduce a minimum threshold τ :

$$\text{Ready} = \begin{cases} 1, & \text{if } qvp \geq \tau \\ 0, & \text{otherwise} \end{cases}$$

6.2. Stakeholder Engagement and Change Management


An integrated framework for enterprise data products—including Agentic AI capabilities—offers extensive benefits for a range of strategic decision-making scenarios in the insurance and retirement industries. A successful implementation roadmap will prioritize stakeholder engagement, emphasizing improved trust and shared ownership of the products.

For products that influence shaping and delivery strategies, multi-dimensional stakeholder engagement is key to maximizing adoption and usage for the benefit of the entire enterprise. Regulatory-oriented enterprise data products provide regulatory stakeholders with the highest level of trust and transparency. Attested provenance for data sources, algorithm training, feature construction, scoring data, network and other settings reinforces the regulatory trust. Risk simulations or education of the regulators can amplify the trust. Providing adequately auditable and transparent systems makes it easier for the enterprises to foster enterprises’ trust. When the regulators’ trust is established using the above techniques, it becomes easier for others to trust



these products. Engaging the auditor stakeholders during design to clearly outline and document explainability features results in an auditability blueprint. Catering to the needs of core product development and pricing development teams during design and implementation builds product delivery teams' trust in the products, which in turn increases enterprise-wide adoption.

Since Agentic AI introduces safety and trust concerns, special care should be taken to tutorialize and educate core business users related to these risks. A training material repository with an easily understandable list of Agentic AI-specific risks and challenges—such as hallucination, misuse, generation-focused responses rather than conclusion-focused responses, differences from traditional chatbots, and less reliable information sources—enables conversion of novice users to experience users and increases overall trust. Providing new users with hands-on formal training and monitoring their responses to potential Agentic AI misuse can ameliorate risks due to misuse.



7. Challenges and Mitigation Strategies

Two fundamental challenges arise when incorporating Agentic AI into Enterprise Data Products for Strategic Decision-Making in the Insurance and Retirement domains: ensuring seamless technical integration across involved systems; and guaranteeing sufficient data quality, availability, and provenance for effective long-range decision-making.

Interoperability between Agentic AI solution components—such as the underlying Infrastructure-as-Code modules, Simulation & Digital Twin Systems, or the State Management Solution—and Enterprise Data Products, which may reside in other technology



or infrastructure silos, is essential for streamlined data transfers and successful detailed information-access requests. Such coordination can often be achieved efficiently using wrapper components, though insufficient technology interoperability may also necessitate costly data movement to and from systems, producing unwanted latency and higher costs that may partly offset the anticipated benefits of Agentic AI adoption.

Furthermore, Agentic AI typically operates in longer-range decision-making environments that require the integration of various data-product outputs to enable trustworthy predictions. If these Agentic AI System Inputs are not delivered with appropriate volume, quality, and time on-target, Machine Learning models may make inappropriate recommendations that, if accepted, could reduce profit, create unacceptable reputational risk, or even lead to insolvency. Continued focus on data-lineage management practices is therefore crucial. Data quality—whether inherent or intrinsic—relates to the characteristics of the data itself, such as accuracy, consistency, and timeliness. Data availability addresses whether sufficient volume, on-target delivery time, and completeness of specified elements exist for a given task.




Fig 4: Challenges in Agentic AI Development

7.1. Technical Interoperability

The diversity and fragmentation of technology platforms and ecosystems deployed across enterprise and partner systems can pose challenges for introduction of agentic AI into enterprise data products and applications. These differences can result in agentic AI failing to recognize objects, services, and interfaces made available to it through its environment, impacting both its operational performance and the quality of the outcome.

Agentic AI developers and implementors should define a comprehensive technology interoperability vision and strategy, detailing optimum interfaces and service specifications across the architecture, production, and operational layers of the enterprise data product or application being implemented and the other enterprise data products or applications with which it will be frequently interacting. Internal enterprise technical standards, ensuring compatibility among enabling technologies with connected agentic AI systems and enterprise data products, should also be considered. These standards should include best practice guidelines covering probable environments of use, data quality, and other factors influencing agentic AI operational



quality. Monitoring agentic AI operational performance, especially in terms of outcome precision and relevancy, through technical performance dashboards provides a method for determining whether technical interoperability issues are arising. If so, suitable interventions should be identified, prioritized, and addressed.

At the same time, it is important to recognize that enterprise data products and applications, by their very design and purpose, will be subject to constant evolution to cater for changing business conditions and requirements. As a result, the creativity, initiative, and adaptability of agentic AI will lead to its successful deployment by enterprise data products and applications throughout an organization's ecosystem, irrespective of the technical challenges faced during its introduction stage.

7.2. Data Quality and Lineage

As enterprise data products become more sophisticated and increasingly put business-critical decisions in the hands of Agentic AI systems, AI-facilitated data tools can help improve the quality of training and operational datasets. Such systems continue to build higher quality datasets for downstream applications, either as side-effects or through their direct design. For example, clustering approaches may curate higher quality training datasets for models to generalize across testing scenarios. Large Language Models (LLMs) hosted as Data Products can automatically classify data provenance at scale, inform data consumers of data sensitivity or provide contextual guides of how data may or may not be used by Downstream Products that interact with such Data Products. Maintaining data lineage entirely remains a difficult, open problem. Sensor fusion methods have an explicit goal of tracing data provenance across transformed datasets, for Descendant Data Products generated through such transformations.

A "data-grounded testing set" can be synthesized, where for each original testing sample, its nearest-neighbor in the original data and the nearest-neighbor in the Downstream Product-space are shown. By performing common-sense reasoning across these samples, it can be verified if falsified scenarios are logically inconsistent, confirming that the discrepancy between falsified openings and closed openings is a property of the Downstream Product that follows naturally from the training set or sum of its parts. Such aids can directly harness Data Infrastructure and the complete history embedded into high-quality kept Data Products and can also build trust and credibility for Data Products in their crossing of Training and Testing domains when the Data Infrastructure becomes rich enough to justify such exploits.

8. Conclusion

The integration of agentic AI into enterprise data products for strategic decision-making in cyber insurance and retirement solutions is described, identifying four aspects in which agentic AI has the potential to support managerial decision-making in these domains: underwriting and risk assessment, product personalization and pricing, product marketing and distribution, and supply chain decision-making. The provision of insurance and retirement solutions relies on complex business architectures and ecosystems involving a plethora of stakeholders. Development is therefore prioritized in partnership with a select group of stakeholders and market leaders. The integration roadmap is designed as a maturity model that spans three phases of increasing complexity, starting with the introduction of agentic AI capabilities into data and analytics platforms and extending to more complex enterprise data and decision products. Technology and data governance challenges that need to be addressed along the journey are identified and discussed.



To navigate the design choices and complexity of these solutions, an agentic-AI-enhanced enterprise data product operating model has been proposed. This extends earlier concept work on industry cloud services for data-driven solutions by encompassing agentic AI capabilities and interactions with solution ecosystems. The enterprise data product operating model acts as a blueprint for solution providers to develop agentic-AI-based solutions in partnership with a data and analytics platform provider and ecosystem stakeholders. The approach enables solution providers to develop the necessary trust, safety, and governance elements to mitigate concerns that executives may have about the introduction of an agentic AI component into enterprise products that they plan to use for making managerial decisions. The use of an agentic-AI-enhanced enterprise data product operating model provides insurance and retirement companies with a framework for integrating agentic AI into enterprise data products for these domains.

8.1. Emerging Trends

A discussion of the trends shaping the further integration of AI into enterprise data products, particularly within strategic decision-making contexts for insurance and retirement solutions.

In 2023, companies interested in growth are looking to leverage the power of AI. Evidence of this broad interest can be seen in the adoption of generative AI products and the willingness of companies to invest in generative AI start-ups. As this genre of AI continues to evolve, the future development of enterprise data products is likely to increasingly incorporate Agentic AI components. Agentic AI can be characterized as the use of AI that provides users or an organization with the ability not only to seek, curate, and discover data but also to act autonomously or semi-autonomously in executing specified operations, such as creating a report on customer behavior or suggesting the next best action for a customer service advisor. The key capabilities that deploy Agentic AI within enterprise data products are the ability to naturally describe tasks, receive contextually relevant data in response, possess topic expertise, code in multiple languages, and perform autonomous self-improvement through continuous learning and auditing.

The incorporation of Agentic AI into enterprise data products is becoming more evident in insurance and retirement-related solutions. For many companies, product lines in these domains are mature; sales and profitability require ongoing investment in product stewardship and operations to maintain sound underwriting performance and ensure relevance to changing customer needs. Data-driven, strategic decision-making for operations using these product lines can benefit from enterprise data products, but these span a diverse range of specialist environments and technical skills. Agentic AI capabilities incorporated into an enterprise data product enable users from a citizen data science and analytics perspective to easily access the underlying data-mart corporate semantic layer to generate and exploit insights and predictions using popular open-source or proprietary technologies in a self-service and low-cost manner.

9. References

- [1] Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv.



- [2] Chen, L., Zaharia, M., & Zou, J. (2021). How is ChatGPT's behavior shaped by reinforcement learning from human feedback? *ACM Queue*, 21(1).
- [3] Davenport, T. H., & Mittal, N. (2022). *All-in on AI: How smart companies win big with artificial intelligence*. Harvard Business Review Press.
- [4] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv*.
- [5] European Insurance and Occupational Pensions Authority. (2021). *Artificial intelligence governance principles: Towards ethical and trustworthy AI in the insurance sector*. EIOPA.
- [6] Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694.
- [7] Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., & Kagal, L. (2018). Explaining explanations: An overview of interpretability of machine learning. *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics (DSAA)*.
- [8] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- [9] Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 1–42.
- [10] Hassan, N., & Jabeen, F. (2022). Explainable AI for insurance risk assessment: A review of interpretable models and applications. *Expert Systems with Applications*, 207, 117900.
- [11] High-Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy AI*. European Commission.
- [12] International Organization for Standardization. (2023). *ISO/IEC 42001:2023 Artificial intelligence—Management system*. ISO.
- [13] International Organization for Standardization. (2022). *ISO/IEC 23894:2022 Artificial intelligence—Risk management*. ISO.
- [14] International Organization for Standardization. (2022). *ISO/IEC 22989:2022 Artificial intelligence—Concepts and terminology*. ISO.



- [15] Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human–AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586.
- [16] Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260.
- [17] Kearns, M., & Roth, A. (2019). *The ethical algorithm: The science of socially aware algorithm design*. Oxford University Press.
- [18] Kroll, J. A. (2021). Out of the black box: The legal and ethical implications of algorithmic decision-making. *Harvard Journal of Law & Technology*, 34(1), 1–43.
- [19] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
- [20] Mikalef, P., Fjørtoft, S. O., & Torvatn, H. Y. (2021). Artificial intelligence in the decision-making context: A systematic review. *Information Systems Frontiers*, 23, 1–25.
- [21] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*.
- [22] Molnar, C. (2022). *Interpretable machine learning (2nd ed.)*. Independently published.
- [23] Nasser, M., & Alizadeh, M. (2023). AI-driven underwriting and claims automation: Opportunities and risks in insurance operations. *The Geneva Papers on Risk and Insurance—Issues and Practice*, 48(3), 501–527.
- [24] National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST.
- [25] National Institute of Standards and Technology. (2023). *AI RMF Playbook*. NIST.
- [26] Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- [27] Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT)*.



- [28] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
- [29] Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- [30] Shapiro, S. S., & Jain, H. (2022). Algorithmic governance and compliance in financial and insurance services: A practical framework. *Journal of Financial Regulation and Compliance*, 30(3), 312–331.
- [31] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
- [32] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2014). Intriguing properties of neural networks. *International Conference on Learning Representations (ICLR)*.
- [33] Tariq, H., & Khan, M. (2021). Predictive analytics and machine learning for retirement planning: A systematic literature review. *Decision Support Systems*, 142, 113474.
- [34] The Geneva Association. (2020). *Artificial intelligence and insurance: Global perspectives on governance and regulation*. The Geneva Association.
- [35] Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- [36] Varian, H. R. (2018). Artificial intelligence, economics, and industrial organization. In A. K. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 399–419). University of Chicago Press.
- [37] Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887.
- [38] World Economic Forum. (2021). *Global AI action alliance: Guidelines for responsible AI adoption*. World Economic Forum.
- [39] Zheng, L., Li, J., & Wang, X. (2024). Agentic LLM systems for enterprise decision support: Architecture patterns and risk controls. *IEEE Software*, 41(2), 45–54.

