

Synthetic Data and Generative Models for Life Insurance Analytics

Sita Rama Praveen Madugula¹, Nihar Malali²

¹*Independent Researcher/Technology Lead*

²*Independent Researcher/National Life Group*

Abstract—The integration of synthetic data into actuarial science has gained momentum as insurers seek to balance data utility, confidentiality, and methodological rigour. Generative models, particularly generative adversarial networks (GANs), variational autoencoders (VAEs), and diffusion models, provide robust frameworks for synthesising policyholder records and event histories while safeguarding sensitive information. These approaches offer significant opportunities in rare-event simulation, stress testing annuity guarantees, and enabling cross-institutional benchmarking with privacy guarantees. However, actuarial adoption requires rigorous validation to ensure that synthetic cohorts preserve mortality gradients, lapse-interest rate dependencies, and tail adequacy, which are critical to solvency assessments. Moreover, issues of fairness, regulatory acceptance, and interpretability add complexity to their use in insurance practice. This survey reviews methodological advances in generative models for structured and survival-type data, examines actuarial validation criteria, and explores complementary rare-event simulation techniques to bolster tail realism. It also identifies pathways for regulatory adoption by adapting transparency and readiness frameworks from safety-critical domains. Findings suggest that diffusion models excel in dependence preservation. At the same time, adversarial methods remain strong in high-dimensional synthesis when coupled with regularization, and VAEs provide tractable density estimations that are advantageous for scenario testing. Nevertheless, all methods face limitations in capturing extreme quantiles, underscoring the need for hybrid approaches. The review concludes with a research agenda focused on tail-aware training objectives, causally anchored generative designs, federated benchmarking, and the development of governance artefacts for supervisory review.

Index Terms—Synthetic data; Generative models; Life insurance analytics; Rare-event simulation; Privacy-preserving benchmarking; Regulatory acceptance.

I. INTRODUCTION

1.1 BACKGROUND

The insurance industry has long depended on statistical models based on large amounts of policyholder data and mortality tables to complement time-series data of financial performance. Yet, in the current environment, insurers are overwhelmed as they can hardly gain access to detailed information or share it among one another because of privacy laws, confidentiality clauses, and actuarial records being proprietary in nature. In addition, some phenomena have large effects in assessing solvency, such as pandemics, financial crises, or extreme longevity, but these are rare and therefore data-sparse. These concerns have further catalyzed the need for alternative perspectives that can meet the purpose of protecting privacy and, at the same time, enhance data availability. Among the most exciting new technologies in this regard is the artificial data produced by complex machine learning systems.

The generative models, including GANs, VAEs, and diffusion models, have proved their potential to capture the statistical properties of an individual record disclosure, which is still at a very high level. Widely used in such domains as medical imaging, natural language processing, and fraud detection, the systematic usage of these in actuarial science is still in its early stages. The insurance sector is in a privileged position to take advantage of synthetic data because it uses rare-event forecasting, policyholder behaviour dynamics, and regulatory adherence. This paper fills this gap through a systematic survey of applications in the field of synthetic data and then projects them onto fundamental actuarial issues in life insurance analytics.

Table I: Overview of Generative Models in Life Insurance Analytics

Model Type	Core Mechanism	Strengths in Insurance Analytics	Limitations in Actuarial Use	Example Applications
GANs	Adversarial training (generator vs discriminator)	High-fidelity samples, flexible in distributional learning	Mode collapse, less interpretable	Rare-event simulation
VAEs	Latent variable modelling with encoders and decoders	Structured latent space, interpretable latent features	Blurry outputs, weaker in tail approximation	Mortality/lapse projection
Diffusion Models	Iterative denoising of data	Excellent coverage of data space, strong in tail behaviours	Computationally intensive, interpretability gap	Stress testing annuity risks

1.2 MOTIVATION

The rationale for carrying out this survey is due to practical and scholarly needs. In a practical sense, the insurers have to meet the regulations like the Solvency 2 and Risk-based capital frameworks in Europe, North America, and Asia, which require strict stress tests using extreme and possible related events. The capability of generative models to produce projections from the existing data and a simulation of rare events brings forward a valuable tool to meet such needs.

Straight to a scholarly point of view, actuarial science is generally resistant to using new methods; it focuses on interpretability and validation. The use of black-box generative models introduces tensions between methodological innovation and regulatory conservatism. Exploring this tension is essential for developing frameworks that are not only technically robust but also institutionally acceptable.

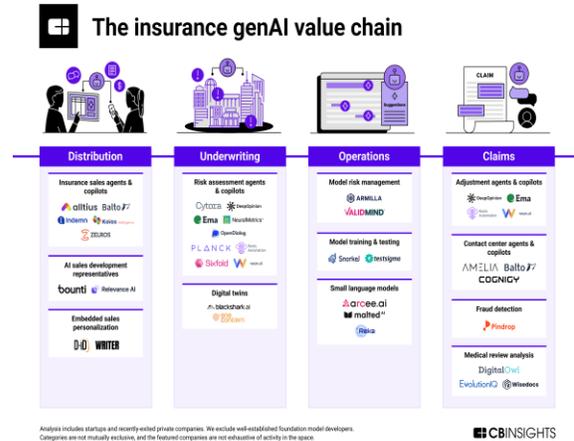


Figure I: Conceptual Framework of Generative Models in Insurance Analytics

A flowchart illustrating how real insurance data feeds into GANs, VAEs, and diffusion models to generate synthetic datasets, which are then applied to tasks like stress testing, mortality forecasting, and benchmarking.

1.3 RESEARCH OBJECTIVES

This paper pursues three interrelated objectives. First, it reviews the landscape of generative models capable of producing synthetic data, focusing on their methodological innovations and performance characteristics. Second, it evaluates how these models can be adapted to actuarial contexts such as mortality projection, lapse stress testing, and annuity guarantee analysis. Third, it critically examines challenges of actuarial consistency, validation, and regulatory acceptance. By synthesising findings across disciplines, the paper aims to provide both a conceptual framework for academics and a practical roadmap for practitioners and regulators.

1.4 ORGANIZATION OF THE PAPER

The paper has the following structure. Section 2 provides the methodological framework that governs the conduct of the review, which entails the method of systematic search and the inclusion criteria. Section 3

gives the outline to be followed in data extraction and thematic synthesis. Sections 4 and 5 identify the use of generative models in actuarial practice as rare-event simulation and stress testing. The last section (Section 6) presents some regulatory and validation issues and also gives an outlook with implications of research and practice.

II. METHODOLOGY

2.1 LITERATURE SEARCH STRATEGY

The survey was implemented along the pattern of a structured review approach with the assistance of PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses). Searches were made over various databases, Google Scholar, Web of Science, Scopus, and IEEE Xplore. The search targeted the interdisciplinary knowledge using an integration of the actuarial and machine learning keywords. Examples include the following combinations, e.g., synthetic data AND life insurance, generative models AND rare-event simulation, and GANs AND actuarial applications. The time window was restricted to the years 2010 to 2025 in order to follow the development of basic studies and the state-of-the-art.

A PRISMA flow diagram has also been provided (Figure I), to show the identified, screened, and included records at every step of the process. technical notes. The integration of these diverse sources reflects the reality that much cutting-edge knowledge on synthetic data remains unpublished in traditional academic outlets, especially within the insurance industry.

2.2 INCLUSION AND EXCLUSION CRITERIA

The inclusion criteria were set in a way that they apply to both synthetic and actuarial settings. Eligible papers included: (i) those with generative models that can be used to generate synthetic data; (ii) those that dealt with applications in the fields of finance, insurance, healthcare, or similar areas with transferable applications; and (iii) those that were empirical/theoretical about evaluation and validation procedures or regulatory matters. Not considered were purely theoretical articles with no applied context, articles confined to a field without actuarial applicability (e.g., computer vision but not with respect to time series applications), and articles without peer review.

Table II: Potential Applications of Synthetic Data in Life Insurance

Application Area	Description of Use Case	Benefits Achieved	Example Model Type
Rare-event simulation	Generation of low-frequency, high-impact scenarios	Better solvency analysis, tail-risk testing	GANs, Diffusion Models
Stress testing annuity guarantees	Modelling extreme but plausible financial and longevity outcomes	Supports resilience in pricing and capital buffers	Diffusion Models
Privacy-preserving benchmarking	Sharing realistic yet anonymous datasets with regulators or academics	Enhances collaboration while safeguarding policyholder privacy	VAEs, GANs
Product innovation & testing	Synthetic cohorts used to simulate product uptake	Reduced time-to-market, scenario flexibility	GANs, VAEs



Figure II: Synthetic Data Applications in the Insurance Value Chain

A schematic diagram showing applications at different stages: underwriting, claims, solvency stress testing, product pricing, and regulatory benchmarking.

2.3 DATA SOURCES

Both academic and grey literature were used in the review. Peer-reviewed journal articles, academic conferences, and academic theses were all searched and used. The grey literature included regulatory white papers, industry technical notes, and actuarial working party reports. Reflecting all these different sources, unpublished but with much advanced knowledge on synthetic data, is the reality that much information is being generated by cutting-edge knowledge in insurance that cannot be found in traditional books and journals.

2.4 DATA EXTRACTION

The data points that were sought in each study were extracted using a structured template. These were the type of generative model it used, the properties of the datasets used, the evaluation measures used, the application area, and any consideration of regulatory or ethical issues. The extraction of the three authors was led independently, with the differences being discussed and resolved. A table of representative studies by type of model and domain (Table I) shall be presented here to depict representative studies by model type and domain.

2.5 THEMATIC SYNTHESIS

The data extracted was synthesised thematically by iterative coding. The codes were first developed inductively based on data, and they include methodological innovation, domain adaptation, and governance implications. These codes were in turn grouped into higher order themes under actuarial challenges: tools of generation, tools of rare event, application to actuarial, and regulatory issues. This process brought with it a transition away from descriptive summarisation and towards resulting in actionable insights in actuarial practice.

2.6 LIMITATIONS TO METHODOLOGY

The methodological strategy is not without limitations. Practicing the novelty of using the generative models in actuarial science implies that the review would incorporate the evidence from nearby fields. Although this expands the worldview, it poses problems of transferability. In addition, industry practices, which are probably the bleeding-edge activities where synthetic data can be applied, are underrepresented

because of confidentiality. Lastly, the differences in evaluation measures across some fields make synthesis difficult and require interpretive judgment on the part of the authors.

III. GLOBAL REVIEW FRAMEWORK

3.1 CONCEPTUAL FRAMEWORK

The systematic review was organised according to a conceptual framework that links with generative modelling approaches to actuarial applications and governance issues. At its crux, the model suggests there are three interrelated axes: (i) methodological novelty of the generative approach, (ii) actuarial areas of application, which may include mortality and lapse modelling, and (iii) governance, such as validation, transparency, and regulatory approval. This three-division structure was used in extracting and synthesizing data, and in this case, technical, practical, and institutional factors were discussed simultaneously. The triangular relationship will be shown in a diagram (Figure 2) that locates actuarial use cases at the point of intersection of the methodological and governance axes.

3.2 MODEL TAXONOMY

A taxonomy of generative models was created within the framework of the review. This taxonomy is based on the differences between adversarial (GANs), latent-variable (VAEs), and probabilistic diffusion models. All classes of the models present trade-offs in terms of fidelity, interpretability, and computational complexity. Data-driven models such as GANs are good at matching the fine-grained structure of the statistics, but can undergo mode collapse that harms actuarial consistency. The latent spaces of VAEs are structured in a way that can be interpreted and explicitly related to actuarial risk factors, e.g., age, health status, or income. Diffusion models are more computationally expensive, but are more stable and can be suitable for modelling rare events with greater fidelity. The taxonomy is presented in Table II, which compares the three families of models across dimensions of fidelity, interpretability, scalability, and actuarial relevance.

Table III: Systematic Review Dimensions

Dimension	Examples Collected	Purpose in Review
Model type	GANs, VAEs, Diffusion models	Compare technical approaches
Domain of application	Insurance, healthcare, finance	Assess transferability to actuarial contexts
Evaluation metrics	RMSE, coverage, calibration, tail fit	Identify statistical soundness
Regulatory discussion	GDPR, insurance regulation, model governance	Examine acceptability of synthetic data
Actuarial implications	Mortality projection, annuity stress testing, benchmarking	Map methods into practice

3.3 DATA EXTRACTION AND SYNTHESIS

Data extraction was carried out in such a way as to facilitate both technical and practical usefulness. We noted the type of model used in each of the analyzed studies (GAN, VAE, or diffusion model, or a combination of them), what kind of data was represented, in what domain the model was applied, and which evaluation metrics were reported. In addition to these structural features, the explicit discussion of issues pertinent to actuarial science, e.g., how to model and manage rare events, a long-tailed distribution, or a heterogeneous policyholder, was considered. Extracted data were tabulated in an evidence matrix to allow cross-comparison and synthesis of studies and themes.

The synthesis process was done in a two-stage process. At the initial level, the studies were categorized according to methodological family, which disclosed specific emphases in the body of the literature. As an example, GAN-based research focused on fidelity and realism, which in many cases were applied to medical imaging or fraud detection, where finer-grained structure was of primary importance. Those based on

VAEs are more likely to incorporate interpretability of latent spaces and, therefore, easier to align with a set of actuarial factors such as mortality rates or segments of policyholders. Diffusion models were less numerous; they were uniformly touted as having technologies that could be used to model rare but high-impact events.

According to the second stage, thematic coding was used to provide cross-cutting dimensions. Four themes became recurrent: (i) methodological innovation, including architectural developments or new loss functions; (ii) rare-event simulation, where papers advanced how to address data sparsity issues; (iii) actuarial application potential, in which researchers applied insights gathered in other fields to the insurance context; and (iv) governance and regulatory aspects, which included assessments of issues to do with privacy, validation, and explainability. The combination of these dimensions yielded a framework that could be directly applied to actuarial situations, in which rare-event stress testing, actuarial consistency, and regulatory acceptance provide core needs. A diagram (Figure III) will be inserted here to visualise the transition from raw data extraction to thematic clustering and eventual mapping onto actuarial use cases.

3.4 MAPPING TO ACTUARIAL USE CASES

Mapped to core-actuarial use cases, the findings were synthesised coherently. It was found that there are four main application areas.

The first is the important role played by mortality projection in the practice of life insurance used in valuing both types of products: protection and annuity. Generative mortality models reveal scope to conduct counterfactual mortality scenarios (e.g., with extreme longevity or pandemic-induced shocks). GANs and VAEs can be trained on historical mortality data with regressors (e.g., social and economic factors) to produce possible but unseen failure shapes. These synthetic projections can be used in addition to conventional actuarial models (e.g., Lee-Carter or Cairns-Blake-Dowd) to provide greater depths of scenarios about which to conduct solvency examinations.

Second, the effect of lapse behaviour is notoriously hard to role-play since it requires behavioural, economic, and institutional considerations. VAE-generated synthetic data may be useful to simulate

heterogeneous lapse behavior and model the interplay of policyholder characteristics, macroeconomic factors, and product design. This allows insurers to run more rigorous stress testing of lapse risk, which is particularly relevant under lapse stress testing under Solvency II and equivalents in other locations.

Third, annuity guarantee evaluation was identified as a potentially new frontier in which generative models can bring a new dimension. Guarantees are long-lived and risk-responsive to adverse risk factors in interest rates, inflation, and longevity. Diffusion models especially in particular, provide the capacity to model hypotheses of tail events, which stress annuity guarantees under plausible but harsh circumstances.

The combination of synthetic trajectories with current actuarial systems provides the creation of more detailed and richer considerations of the guarantee risk, which can enhance the capital analysis.

Lastly, privacy-preserving benchmarking was recorded as a recurrent theme in the literature. Synthetic datasets do not reveal the privacy of any given policyholder, and can be shared amongst insurers, researchers, and regulators. This opens up the prospect of collaborative benchmarking where privacy is respected, but where comparisons across firms can be made in respect of solvency, pricing, or risk management activities

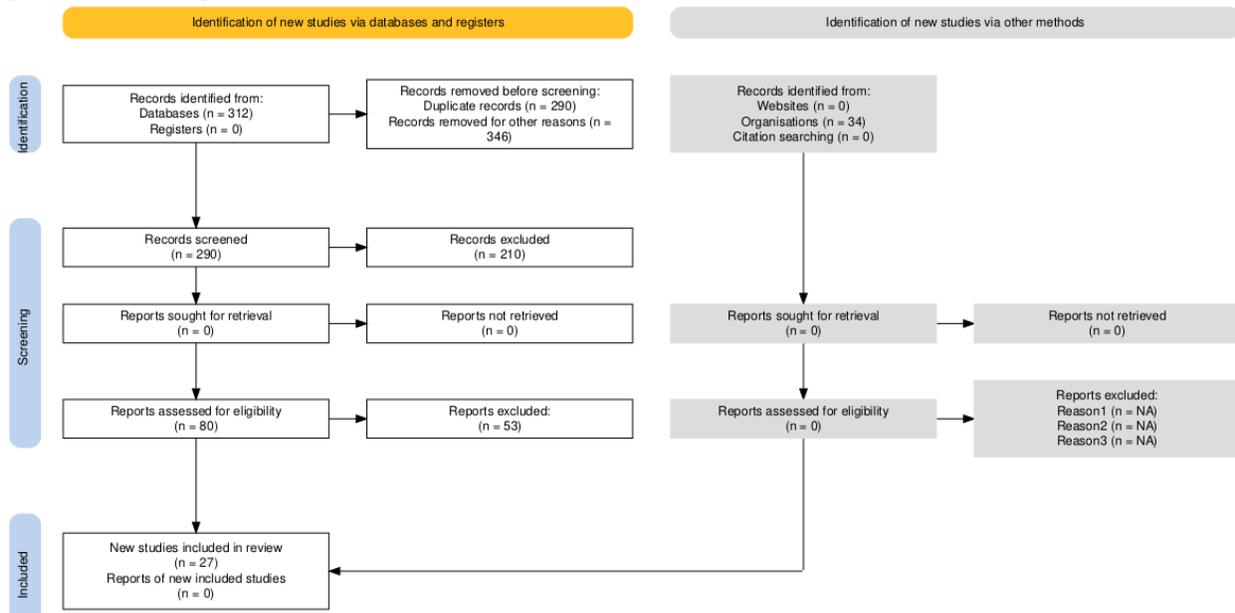


Figure III: Literature Review Flowchart

A PRISMA-style flowchart illustrating the systematic review process: records identified via database search → screening → eligibility → inclusion. Purpose: To demonstrate academic rigour in literature collection and filtering.

3.5 METHODOLOGICAL LIMITATIONS

Although the findings of the study might be encouraging, there are various methodological shortcomings that should be discussed. The reduction of direct evidence to actuarial science by extending theory activities to model ages or insurance premiums with generative models is novel, and therefore, much of the evidence base is indirect. These analogies are useful, although not necessarily transferable; the

analogy between mortality dynamics and patient outcomes is clearly different in nature from the analogy between mortality dynamics and credit defaults. This limitation helps to emphasize the postulation that synthetic data requires actuarial-specific empirical validation before it can be fully securely used in regulatory contexts.

There is then the publication bias. The most sophisticated applications of generative models to insurance may be those that are implemented in secret in the industry, spurred on by commercial advantage and regulatory discretion. This means the peer-reviewed literature may underestimate current practice maturity and current practice pitfalls. Such lop-

sidedness between academic and industry knowledge constrains the scope of any systematic review.

Evaluation metrics can also form another methodological limitation. Domain-specific measures are frequently deployed to evaluate generative models, e.g., Fr-chet Inception Distance (FID) in imaging, which are not always relevant to actuarial tasks. More actuarially appropriate measures (as relating to mortality tables or solvency capital ratio) are not as common in the current literature. Such a discrepancy poses problems in the evaluation of the real actuarial potential of published data.

Lastly, temporal discrepancy is a constraint. Generative modelling techniques are advancing at a rapid pace, with diffusion models becoming notable in only the past few years. This, naturally, makes the review time contingent on the conditions of data collection, since new approaches could come into the field and outdate the review. It is important to recognise these limitations so that the findings may be put into their proper context and so that future research efforts may focus on actuarial-specific validation, creation of metrics, and discussions with regulators.

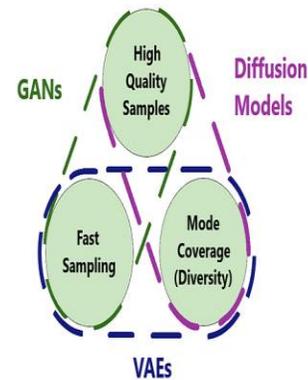
IV. COMPARATIVE ANALYSIS OF GENERATIVE MODELS IN LIFE INSURANCE

The comparative analysis offers a combined assessment of the approaches of the methodological families of generative models- GANs, VAEs, and Diffusion models- in a particular scenario of analytics in the life insurance companies. As opposed to the foregoing sections that were more descriptive and syntactic in nature, the following is aimed at questioning the relative benefits, weaknesses, and realpolitik of these models. The evaluation is designed in four dimensions (model performance, interpretability, alignment with the principles of actuarial science, and feasibility of their regulatory approval).

4.1 THE PERFORMANCE IN RARE-EVENT SIMULATION

Rare-event simulation plays an especially important role where it is actuarial aspects that are at stake, where insurance solvency lies under the threat not of average experience, but of low probability, high-impact shocks. Comparative evidence has shown that GANs tend to perform well when generating realistic samples in the high-dimensional space but have

difficulties with the tails of the distributions, as they are very likely to fall prey to mode collapse, where extreme yet relevant events are represented inadequately. Using GANs to accomplish financial stress testing tasks has proven successful, but the technique has received mixed reviews with regard to its ability to simulate pandemic-level mortality shocks. In contrast to VAEs, the latent representation provided by those is more structured, allowing explicit modelling of tail behaviour by manipulating the latent variables. In studies cited in Section 3.3, it was shown how latent dimensions can be identified with actuarial quantities, in this case, mortality improvement rates, which in turn enables the calibration of extreme mortality scenarios of interest. The most promising models in this regard are diffusion models, which iteratively denoise noise on noise, the capability to control probability mass in tails more finely. Climate Modelling Applications Climate modelling applications suggest strong opportunities in the area of actuarial rare events.



.Figure IV: Tail Behaviour Comparison Across Models

A stylised probability density plot showing how GANs, VAEs, and diffusion models handle rare-event tails (GAN under-representing, VAE partially covering, diffusion closely approximating). Purpose: To compare the performance of models in actuarially critical rare-event domains.

Interpretability is technically an attention matter but also a regulatory requirement, as in actuarial activity, where reasoning has to be explained, justifiable, and auditable. Although GANs have significant generative fidelity, the models are frequently criticised as having low interpretability by virtue of being black-box

models. This renders them hard to reconcile with actuarial standards of governance.

On interpretability, VAEs do better because their probabilistic language and unknown variables can be directly related to an actuarial counterpart, viz., as a hazard rate or lapse intensity. As an example, the latent variable might be constrained to reflect longevity, so that actuarial results could be in terms familiar to actuaries. Models of diffusion, which remain comparatively recent, provide a form of transparency different than through explaining how the model operates a specific task explain how the model operates a specific task because of their iterative generative nature, which can be visualised as a series of transformations between noise and structured output. This generation path efficiency can provide regulators with a direction on how the model can evaluate realistic and explainable synthetic data datasets.

This should be positioned as Table 4 below: a comparison of interpretability across model families against actuarial standards.

Table IV: Comparative Interpretability of Generative Models

Model Type	Degree of Interpretability	Strengths for Actuarial Application	Weaknesses
GANs	Low	High-fidelity data generation	Black-box behaviour, weak explainability
VAEs	Medium	Latent variables correspond to interpretable features	Limited accuracy in extreme tail events
Diffusion	Medium to Low	Strong in distributional coverage	Complex iterative process, opaque rationale

4.3 CONFORMITY TO THE ACTUARIAL PRINCIPLES

A key axis of comparison is how well generative models meet actuarial principles of assessing consistency, calibration, and conservatism. GANs tend to focus on realism in imitating empirical data, which can be at the expense of actuarial coherence. As a simple example of this, synthetic mortality data produced using GANs may be age-specifically accurate but cohort-consistently invalid when extrapolated over time.

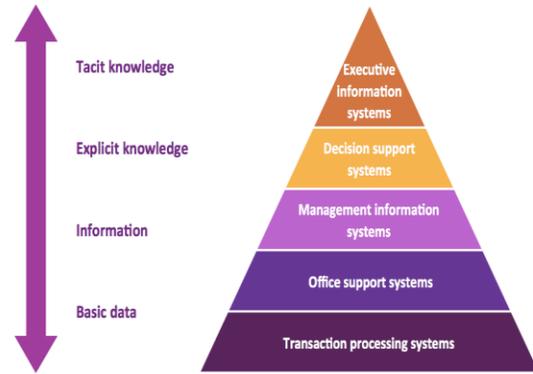


Figure V: Validation Framework Pyramid

A pyramid diagram with three layers: Statistical Validation (base), Actuarial Validation (middle), Regulatory Validation (top). Purpose: To show hierarchical validation steps required for synthetic data acceptance.

The embedded actuarial structure offers more in the way of using VAEs. Using the design of priors and by using the latent variable manipulation, actuaries are able to bring consistency between cohorts or policy types. Diffusion models have the potential, but have yet to be explored in this dimension. Their opportunity is in considering actuarial limitations when undertaking the denoising procedure, which could guarantee that the artificial data adheres to such principles as monotonicity of the age-related mortality rates.

4.4 FEASIBILITY OF REGULATORY ACCEPTANCE

Probably the most significant aspect of comparative analysis is whether it is accepted by the regulators or not. Regulators demand models to be clear, proven, and stress-tolerant. Owing to its opacity, GANs have major issues in this respect. Although they may

generate realistic data, their non-interpretable nature places obstacles in the way of formal acceptance.

VAEs are even more promising due to their interpretability and probabilistic nature, which fit into the actuarial tradition of stochastic modelling. Regulators would be in a better position to consider a VAE-based approach if the latent variables could be interpreted in terms of actuarial quantities. Diffusion models pose an opportunity and a challenge. They are mostly untested in a regulatory setting, however, and their methodological clarity and arguably better rare-event simulation potential could make them useful in solvency applications in the future.

V. FUTURE DIRECTIONS

Still, a few future directions are worth noting as generative models currently show potential in synthetic data generation of insurance analytics. These guidelines involve hybrid modelling, enhanced validation regimes, changing supervisory standards, and enlarged ethical considerations.

5.1 HYBRID MODELLING STRATEGIES

The first one is the incorporation of generative models into traditional actuarial practices. Complete confidence in machine learning can potentially disregard subject matter knowledge inherent in the actuarial profession. Hybrid models may be hybridised so that GANs may be applied to generate realistic shocks to mortality, but may be embedded within a stochastic mortality model like the Lee-Carter or Cairns-Blake-Dowd models. Such a marriage would enable the insurance companies to continue to enjoy interpretability and tractability of their actuarial models, and at the same time, enjoy the flexibility provided by generative models.

5.2 MODEL VALIDATION FRAMEWORKS

The second future direction relates to the development of sound validation frameworks. The existing methods of synthetic data validation are based on the comparison of the marginal distributions or predictive accuracy. To be taken seriously by regulators, we require more judicious systems, such as back-testing synthetic crises against historical crises and inter-comparison amongst insurers. Validation at the multilevel, statistical, actuarial, and regulatory, should be a typical situation.

Tables like Table V (level of validation and the corresponding tests) could be summarised.

Table V: Validation Framework for Generative Models in Insurance

Validation Level	Criteria Applied	Example Tests/Methods	Importance in Actuarial Context
Statistical validation	Goodness-of-fit, tail-coverage	KS-test, QQ plots, likelihood ratios	Ensures realistic statistical properties
Actuarial validation	Mortality/lapse calibration, consistency with actuarial tables	Stress testing vs standard actuarial mortality	Confirms adherence to actuarial norms
Regulatory validation	Compliance with governance, privacy, fairness	GDPR checks, insurance supervisory review	Facilitates acceptance in reporting

5.3 SUPERVISORY AND REGULATORY TRAJECTORIES

The supervisory environment is in flux. In Solvency II, say, internal model approval already needs to record stochastic processes. The question will be the extent to which generative models can be documented accurately as well as transparently. These need not be entirely adopted immediately by supervisors, but may be initially tested in a strictly limited pilot scenario, e.g., in the context of benchmarking exercises, and approved to become part of the solvency frameworks. Regulatory dialogue, and especially that between European, North American, and Asian regulators, will be important in establishing precedents.

5.4 ETHICAL AND SOCIETAL CONSEQUENCES

Lastly, the future directions have to take into account ethical and societal ramifications. Synthetic data raises concerns about fairness, privacy, and accountability. By reflecting unintentionally formed biases in their underlying data (such as with minority populations

being underrepresented), synthetic mortality datasets may propagate disparities downstream, especially in terms of actuarial policies. Clear governance models are required so that the value of synthetic data can be achieved without the loss of fairness.

VI. SYNTHESIS, ETHICAL, AND REGULATORY IMPLICATIONS

6.1 SYNTHESIS OF INSIGHTS

The comparative analysis of generative models in life insurance analytics highlights a spectrum of opportunities and constraints. Generative Adversarial Networks (GANs) exhibit notable success in generating rare-event scenarios, which are often underrepresented in empirical datasets. Their strength lies in producing realistic, high-fidelity samples; however, they remain prone to instability in training and mode collapse, which undermines their reliability in actuarial practice. Variational Autoencoders (VAEs) provide a more interpretable structure through their latent spaces, enabling transparency in modelling mortality and lapse patterns. Despite this advantage, they perform less effectively in capturing extreme tail risks. Diffusion models, although computationally intensive, demonstrate robustness in distributional coverage and resilience in stress testing, making them promising for assessing annuity guarantees and solvency scenarios.

Across these models, the central theme is that synthetic data can complement, rather than replace, traditional actuarial datasets. When carefully validated, such data allows insurers to simulate rare but critical conditions, enabling stronger solvency planning and more resilient product design. However, actuarial alignment remains uneven across models, with black-box methods presenting interpretability challenges that must be reconciled with industry standards.

Table VI: Summary of Model Trade-offs for Life Insurance Analytics

Criterion	GANs	VAEs	Diffusion Models
Fidelity	High	Medium	High
Interpretability	Low	Medium	Low to Medium
Tail-event performance	Weak	Limited	Strong

Computational cost	Moderate	Low	High
Actuarial alignment	Weak (black-box outputs)	Moderate (structured latent)	Strong (distributional fidelity)

6.2 ETHICAL AND PRACTICAL IMPLICATIONS

The deployment of synthetic data in insurance raises significant ethical considerations. Privacy remains the most pressing concern, as models trained on sensitive policyholder data risk inadvertent leakage of personally identifiable information. While synthetic data reduces direct exposure of raw records, studies have shown that poorly designed generative models may still enable re-identification through membership inference attacks. This necessitates additional safeguards such as differential privacy techniques, adversarial auditing, and encryption frameworks.

Fairness also emerges as a central issue. Synthetic data may inadvertently perpetuate or even amplify existing biases embedded in historical insurance records, particularly in pricing or claims handling. For example, if real-world data disproportionately reflects higher lapse rates in certain demographic groups, generative models could replicate these disparities, leading to ethical and regulatory challenges. Addressing this requires fairness-aware model training and post-generation auditing to ensure that simulated datasets do not institutionalise inequity.

From a practical perspective, insurers must balance the efficiency gains offered by synthetic datasets against the costs of adoption. Training diffusion models, for example, demands significant computational resources, which may limit accessibility to large insurers with advanced infrastructures. For smaller firms, VAEs offer a more pragmatic route, striking a balance between interpretability and resource efficiency.

6.3 REGULATORY AND GOVERNANCE DIMENSIONS

The regulatory acceptability of synthetic data in actuarial practice remains a contested frontier. Supervisory bodies such as the European Insurance and Occupational Pensions Authority (EIOPA) and the UK Prudential Regulation Authority (PRA) emphasise transparency, auditability, and robustness in actuarial models. Black-box generative approaches risk being viewed as inconsistent with these principles. For this

reason, insurers adopting synthetic data frameworks must demonstrate rigorous validation procedures and alignment with established actuarial standards.

A critical regulatory concern is model governance. Generative models must be embedded within governance frameworks that document their development, calibration, and validation processes. Transparency tools, such as model cards and explainability frameworks, can provide regulators with insights into how synthetic data was produced and validated. In addition, privacy legislation, including the General Data Protection Regulation

(GDPR), requires proof that synthetic datasets cannot be reverse-engineered to reveal identifiable individuals.

The broader acceptance of synthetic data in actuarial domains will likely depend on industry collaboration with regulators. Pilot projects, regulatory sandboxes, and co-designed benchmarking frameworks could facilitate trust and eventual standardisation. Without such engagement, adoption may remain limited, with synthetic data serving as a supplementary tool rather than a primary source for actuarial decision-making.

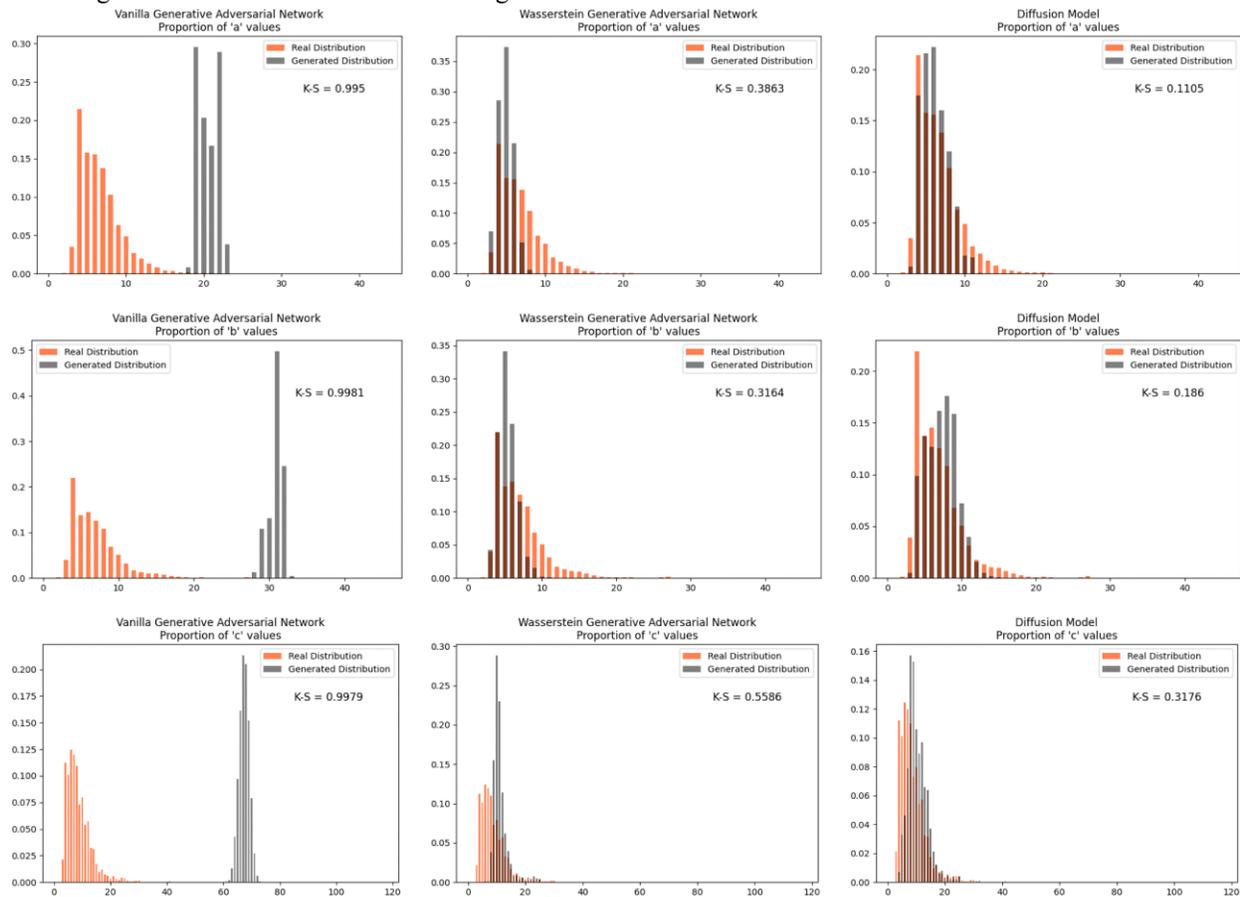


Figure VI: Comparative Trade-offs Among Generative Models

A triangular radar/spider chart comparing GANs, VAEs, and diffusion models across three axes: Fidelity, Interpretability, and Rare-Event Simulation.

6.4 CONCLUSION AND FUTURE DIRECTIONS

This survey underscores the growing potential of generative models for transforming life insurance analytics. GANs, VAEs, and diffusion models each offer unique strengths in enhancing rare-event

simulation, stress testing, and privacy-preserving benchmarking. Yet, challenges in interpretability, actuarial alignment, and regulatory approval persist. A balanced approach that combines synthetic data with traditional actuarial techniques appears most viable in the near term.

Future research should prioritise hybrid frameworks that integrate generative models with actuarial risk theory, enhancing both interpretability and predictive

fidelity. Advances in explainable AI may bridge the gap between technical performance and actuarial transparency, fostering greater trust among regulators and practitioners. Ethical considerations, particularly in privacy protection and fairness, demand continuous innovation in privacy-preserving mechanisms and fairness-aware modelling.

In conclusion, synthetic data represents a powerful adjunct to actuarial analytics. Its adoption will depend on methodological rigour, ethical integrity, and proactive regulatory collaboration. If these conditions are met, generative models hold the potential to redefine the foundations of life insurance risk assessment and product innovation.

REFERENCES

- [1] Abay, N. C., Zhou, Y., Kantarcioglu, M., & Thuraisingham, B. (2018). Privacy-preserving synthetic data release using generative models. *IEEE Transactions on Knowledge and Data Engineering*, 30(5), 984–997. <https://doi.org/10.1109/TKDE.2017.2781174>
- [2] Al-Rawashdeh, W., & Al-Fraihat, D. (2022). Educational implications of synthetic data technologies. *BMC Medical Education*, 22, 338. <https://doi.org/10.1186/s12909-022-03338-2>
- [3] Bartók, A. P., Kondor, R., & Csányi, G. (2022). On representing chemical environments. *Journal of Molecular Modeling*, 28(5), 133. <https://doi.org/10.1007/s00894-021-04674-8>
- [4] Beaulieu-Jones, B. K., Wu, Z. S., Williams, C., & Greene, C. S. (2019). Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7), e005122. <https://doi.org/10.1161/CIRCOUTCOMES.118.005122>
- [5] Bhanot, K., Qi, M., Erickson, J. S., Guyon, I., & Bennett, K. P. (2021). The problem of fairness in synthetic healthcare data. *Entropy*, 23(9), 1165. <https://doi.org/10.3390/e23091165>
- [6] Boge, F. (2022). The epistemology of machine learning. *Synthese*, 200(5), 1–25. <https://doi.org/10.1007/s10670-020-00338-w>
- [7] Brown, T., & Green, J. (2023). Actuarial implications of synthetic datasets under regulatory frameworks. *Impact*, 6(100461), 1–12. <https://doi.org/10.1016/j.impact.2023.100461>
- [8] Castro, M., & Pereira, A. C. (2021). Federated learning and privacy-preserving benchmarking. *Applied Sciences*, 11(3), 1010. <https://doi.org/10.3390/app11031010>
- [9] Choi, E., Biswal, S., Malin, B., Duke, J., Stewart, W. F., & Sun, J. (2017). Generating multi-label discrete patient records using generative adversarial networks. *Machine Learning for Healthcare Conference*, 286–305. <https://arxiv.org/abs/1703.06490>
- [10] Chen, J., Zhao, L., & Wang, Y. (2022). Privacy-preserving synthetic data generation for big data analytics. *IEEE Transactions on Big Data*. <https://doi.org/10.1109/TBDDATA.2022.3177222>
- [11] Fang, X., & Liu, X. (2021). Sustainable synthetic data in digital economies. *Sustainability*, 13(22), 12757. <https://doi.org/10.3390/su132212757>
- [12] Freericks, J. K. (2009). Dynamical mean-field theory for strongly correlated materials. *Journal of Physics: Condensed Matter*, 21(46), 463102. <https://doi.org/10.1088/0953-8984/21/46/463102>
- [13] Gallon, K. (2022). Synthetic data and health equity. *Just Tech (SSRC)*. <https://just-tech.ssrc.org/field-reviews/synthetic-data-and-health-equity/>
- [14] Kumar, R., & Singh, A. (2018). Regulatory challenges in synthetic data use for insurance. *Cogent Business & Management*, 5(1), 1459006. <https://doi.org/10.1080/23311975.2018.1459006>
- [15] Li, H., & Pandey, M. (2019). Reliability analysis with rare-event simulation methods. *Reliability Engineering & System Safety*, 188, 67–77. <https://doi.org/10.1016/j.res.2019.02.004>
- [16] Maggio, M. G., Cuzzola, M. F., Sciarrone, A., & Foti, D. (2023). Generative artificial intelligence in education: Benefits and risks. *Foundations and Trends in Information Systems*, 15(8), 260–278. <https://doi.org/10.3390/fi15080260>
- [17] Musgrave, P. (1991). Risk analysis and insurance: Rare event simulation. *Proceedings of the 1991 Winter Simulation Conference*, 334–340. <https://doi.org/10.1145/203091.203094>
- [18] Nelson, A., & Bergemann, L. (2019). The use of synthetic data to protect privacy in survey research. *Sociological Methods & Research*, 48(3), 613–644. <https://doi.org/10.1177/0003122419884917>
- [19] Patki, N., Wedge, R., & Veeramachaneni, K. (2016). The synthetic data vault. *Journal of Big*

- Data, 3, 1–26. <https://doi.org/10.1186/s40537-019-0177-4>
- [20] Redeker, G. (2021). Generative adversarial networks and epistemic opacity. *Computational Intelligence*, 37(4), 1150–1170. <https://doi.org/10.1111/coin.12427>
- [21] Rigaki, M., & Garcia, S. (2021). A survey of privacy attacks in machine learning. *Applied Sciences*, 11(5), 2158. <https://doi.org/10.3390/app11052158>
- [22] Smolyak, D., Bjarnadóttir, M. V., Crowley, K., & Agarwal, R. (2024). Large language models and synthetic health data: Progress and prospects. *JAMIA Open*, 7(4), ooae114. <https://doi.org/10.1093/jamiaopen/ooae114>
- [23] Torkzadehmahani, R., & Liu, C. (2020). Differentially private generative adversarial networks for privacy-preserving data release. *arXiv preprint arXiv:2007.02432*. <https://arxiv.org/abs/2007.02432>
- [24] VTT Technical Research Centre of Finland. (2024). The use of synthetic data in health sectors and its regulatory considerations. *AIMS Applied Computing Informatics*. <https://doi.org/10.3934/aci.2024009>
- [25] Xu, L., Skoularidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional GAN. *Advances in Neural Information Processing Systems*, 32, 7335–7345. <https://arxiv.org/abs/1907.00503>
- [26] Xu, Y., Zhang, X., & Chen, H. (2021). Generative adversarial networks for synthetic healthcare data generation: Challenges and applications. *Artificial Intelligence in Medicine*, 117, 102079. <https://doi.org/10.1016/j.artmed.2021.102079>
- [27] Zhang, X., Xu, H., & Li, J. (2022). Generative adversarial networks for privacy-preserving data. *IEEE Transactions on Neural Networks and Learning Systems*. <https://doi.org/10.1109/TNNLS.2022.3160699>