

46th India Fellowship Seminar

Date: 11th & 12th June 2026

Life Insurance Technical Case Study on Governance for AI-Based Models

Presented By : Group 2

- 1. Khushboo Asopa**
- 2. Srishti**
- 3. Lakhan Pal Mongia**
- 4. Venkatesh Ronanki**



Case Study - Governance for AI-Based Models

Business context & Issues



A life insurer introduces an AI underwriting and pricing engine using:

- external data
- machine learning
- and customer behavioural information

The system improves profitability and speed.

However,

- Regulators question fairness
- Customers complain about premium differences
- The Board cannot explain the model
- And independent validation identifies possible demographic bias.

Case Study - Governance for AI-Based Models

Agenda



The Board asks the Appointed Actuary to make a presentation covering the following aspects:

- How should fairness and bias be tested in AI-based underwriting and pricing models?
- What human oversight controls are necessary for AI-driven decisions?
- What role should the Board play in AI governance and oversight and to establish AI Ethics?
- Should the model be suspended pending remediation? What factors should be considered?
- What is the evolving role of the actuary in governance of AI-based decision systems?

References

Key Regulatory and Professional requirements



- Insurance Act, 1938
- IRDAI Act, 1999
- Actuaries Act, 2006
- Professional Code of Conduct v4.0
- APS 1 | APS 2 | APS 3 | APS 7 | APS 9 | APS 34
- IRDAI (Actuarial, Finance and Investment Functions of Insurers) Regulations, 2024
- IRDAI (Corporate Governance for Insurers) Regulations, 2024
- IRDAI (Insurance Products) Regulations, 2024
- Master Circular on Corporate Governance for Insurers, 2024
- Master Circular on Life Insurance Products, 2024
- IRDAI (Protection of Policyholders' Interests) Regulations, 2024
- EU AI Act, 2024
- NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers, 2023

Case Study - Governance for AI-Based Models

Agenda



The Board asks the Appointed Actuary to make a presentation covering the following aspects:

- **How should fairness and bias be tested in AI-based underwriting and pricing models?**
- What human oversight controls are necessary for AI-driven decisions?
- What role should the Board play in AI governance and oversight and to establish AI Ethics?
- Should the model be suspended pending remediation? What factors should be considered?
- What is the evolving role of the actuary in governance of AI-based decision systems?

Testing Fairness and Bias in AI Models

Context



Core Purpose of Underwriting:

- **Risk Assessment and Selection** - Evaluate applicant risk profiles accurately through data.
- **Preventing Adverse Selection** - Minimise asymmetric risk from high-hazard applicants.
- **Portfolio Risk Management** - Maintain a balanced and stable book of business.

Core Purpose of Pricing:

- **Financial Solvency and Adequacy** - Ensure premium levels cover future claims and operational expenses.
- **Equity and Fairness** - Charge premiums proportional to the actual risk transferred.
- **Maintain market competitiveness** - Maintain viable, attractive rates within the competitive landscape.

Testing Fairness and Bias in AI Models

The Shift Driven by AI in Underwriting and Pricing Models



Straight-Through Processing (STP)

- **Instant Data Retrieval:** Pulls digital health records and prescription histories in real time.
- **Rapid Approvals:** Onboards low-risk applicants in minutes without invasive medical exams or fluids.

Enhanced Risk Granularity

- **Non-Linear Insights:** Identifies complex risk correlations across massive datasets that manual underwriting may miss.
- **Precise Cohorts:** Establishes highly accurate risk segments, moving past rigid age and gender brackets.

Consistency & Objectivity

- **Uniform Logic:** Applies identical evaluation rules across 100% of standard applications.

Testing Fairness and Bias in AI Models

The Shift Driven by AI in Underwriting and Pricing Models



Consistency & Objectivity

- **Zero Leakage:** Minimises subjective human biases, underwriter fatigue, and processing errors.

Hyper-Personalised Rates

- **Reflective Premiums:** Matches individual premiums directly to true, specific loss probabilities.
- **Fairer Pooling:** Eliminates broad subsidisation of high-risk portfolios by low-risk customers.

Predictive Mortality Modeling

- **Trend Projection:** Leverages machine learning to forecast long-term claims based on evolving health data.
- **Solvency Protection:** Adapts pricing dynamically to shield capital from unexpected macroeconomic shifts.

Testing Fairness and Bias in AI Models

The Risk Landscape: How AI introduce Unfairness and Bias



A. Training Data Reflects the Past (Data Skew)

- **The Mechanism:** AI models learn entirely from historical underwriting files and socioeconomic data.
- **The Unfair Outcome:** AI cannot distinguish between a legitimate risk factor and a consequence of historical inequality, such as lower life expectancy due to unequal socioeconomic conditions.

B. Proxy Variables Carry Hidden Signals (Indirect Discrimination)

- **The Mechanism:** Even when certain fields like race or income are excluded to avoid direct discrimination, the AI might analyse other data points as proxy for these.
- **The Unfair Outcome:** Variables such as PIN code acts as a mirror for race or income. The AI uses them to create discriminatory outcomes that are identical to using direct demographic data.

Testing Fairness and Bias in AI Models

The Risk Landscape: How AI introduce Unfairness and Bias



C. Behavioural Data Penalises the Underserved (The Infrastructure Gap)

- **The Mechanism:** AI relies on real-time data inputs like credit activity or wearable device usage to segment risk.
- **The Unfair Outcome:** Absence of such data, due to limited access to financial means, is misinterpreted as a high-risk indicator, unfairly penalising applicants.

D. Correlation Without Causation (The Black-Box Trap)

- **The Mechanism:** AI finds complex patterns across massive datasets without understanding real logic such as people driving a hybrid car have 15% fewer heart attacks.
- **The Unfair Outcome:** The real reason is that people who buy such cars usually have higher incomes and better access to healthcare. AI has locked onto a random coincidence, without any logical explanation.

Testing Fairness and Bias in AI Models

Measures



A. Input-Level Fairness

- **Proxy Discrimination Risk/Historical Inequality:** Any model inherits all biases present in its training dataset. External and behavioural data frequently carry embedded societal, economic, and regional disparities.
- **Upstream Data Auditing/Proactive Screening:** Shift focus from output metrics to actively cleaning, filtering, and auditing input datasets for historical skew.
- **Independent Assurance:** Hire external, third-party consultants to conduct regular objective reviews of codes and datasets.

Testing Fairness and Bias in AI Models

Measures



B. AI Parameter Testing

- **Test Variable Relevance:** Audit every factor the AI uses to ensure it has a proven, logical connection to actual mortality/morbidity risk.
- **Execute Proxy Detection:** Identify neutral factors (like PIN codes or jobs) that act as proxy discriminators. Further, test what happens when we remove them to see if it makes pricing fairer across different groups
- **Verify Weight Stability:** Test if the weights the AI assigns to factors stays the same across different groups such as locations and wealth levels. This ensures the AI isn't changing its calculation based on hidden bias.

Testing Fairness and Bias in AI Models

Measures



C. Portfolio Benchmarking

- **Baseline Comparison:** Verify AI model outcomes against traditional actuarial models at the sample portfolio level.
- **Shadow Testing:** Process new data through old systems to ensure live AI outputs remain valid and within risk tolerances
- **Variance Tolerances:** Verify and document that all premium and underwriting deviations remain strictly within predefined risk limits.
- **Threshold Investigations:** Isolate and investigate any individual policy where the premium variance exceeds the established threshold.

Testing Fairness and Bias in AI Models

Measures



D. Continuous Monitoring & Automated Calibration

- **Embed Automated Bias Triggers:** Build real-time algorithmic alerts into production environments that flag data anomalies or unexpected demographic spikes immediately.
- **Ingest Continuous Control Data:** Feed localised internal company metrics continuously into tracking dashboards to act as a rolling baseline against external data skews.
- **Audit Profit Variance Anomalies:** Mandate instant algorithmic deep-dives into any sudden, non-linear profitability surges to verify they are driven by genuine risk accuracy rather than hidden demographic exploitation.

Testing Fairness and Bias in AI Models

Measures



E. Complaint Redressal & Model Correction

- **Analyse Policyholder Complaints:** Evaluate all customer complaints systematically to detect and diagnose underlying model biases.
- **Fast-Track Corrective Protocols:** Deploy expedited approval protocols to rapidly implement model rectification and customer redressals.

Case Study - Governance for AI-Based Models

Agenda



The Board asks the Appointed Actuary to make a presentation covering the following aspects:

- How should fairness and bias be tested in AI-based underwriting and pricing models?
- **What human oversight controls are necessary for AI-driven decisions?**
- What role should the Board play in AI governance and oversight and to establish AI Ethics?
- Should the model be suspended pending remediation? What factors should be considered?
- What is the evolving role of the actuary in governance of AI-based decision systems?

Human oversight controls

The need for human oversight controls

Judgement is irreplaceable

Professional judgement remains necessary to assess the appropriateness of data sources and model-driven decisions.

Accountability stays

Legal and regulatory accountability rests with the insurer, the algorithm has no standing.

Models degrade

Model relevance may diminish over time, as the relationships learned from historical data may not remain valid in a changing environment.

Human oversight controls

Risk to control mapping



Risk	Control
Opacity Risk - model decisions cannot be understood or explained	Explainability controls at model and individual decision level
Bias Risk - demographic or proxy variable discrimination	Independent validation with demographic testing and fairness metrics
Accountability Risk - responsibility diffused across algorithm and humans	Respective function owner sign-off authority and designated override professionals

Human oversight controls

Risk to control mapping



Risk	Control
Automation Risk - human review becomes procedural	Human-in-the-loop thresholds with structured escalation tiers
Drift Risk - model becomes less accurate as population or environment changes	Ongoing monitoring with pre-defined recalibration triggers
Regulatory Risk - inability to demonstrate governance to regulators	Model documentation standards and oversight trail feeding into regulatory submissions

Human oversight controls

Measures



Explainability & Transparency

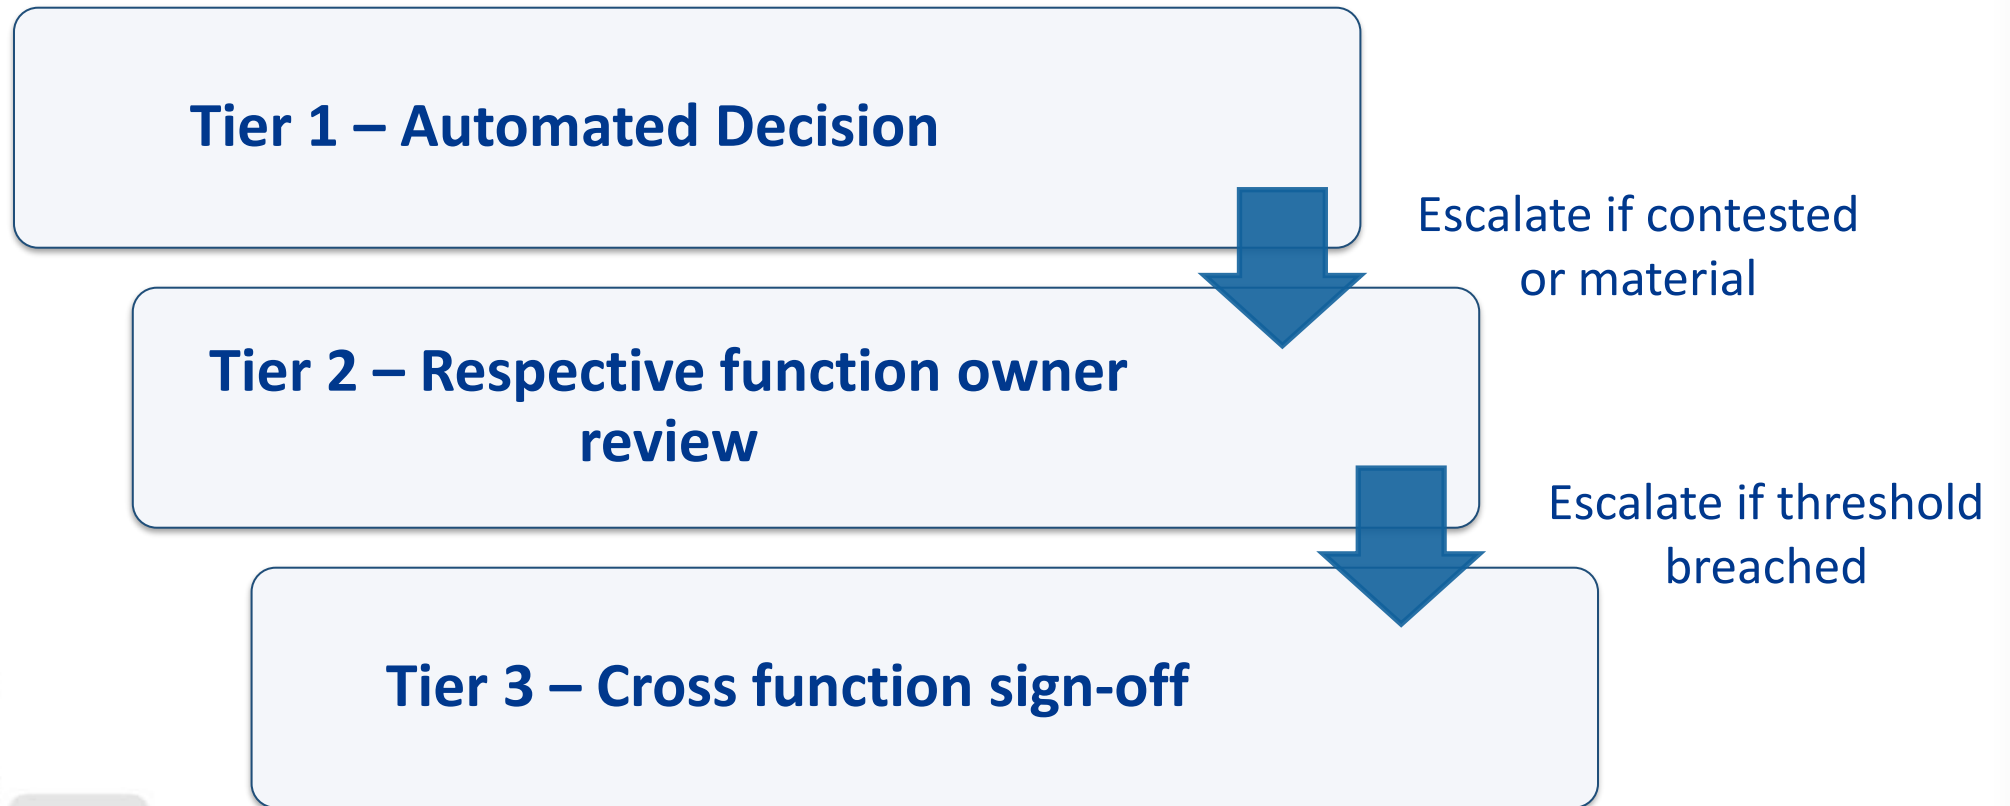
- An unexplainable decision cannot be defended, particularly in a regulatory or legal challenge.
- Assess whether the key drivers of AI-generated decisions can be understood and communicated to the Board, regulators, etc.
- Assess whether the proxy variables used from the external data are appropriate, complete and representative of the insurer's market.

Human oversight controls

Measures



Human-in-the-loop review : Tiered Oversight Framework



Human oversight controls

Measures



Independent Validation

- Peer review of the methodology, data, fairness, robustness, and predictive performance at deployment, periodically thereafter, and upon any material model change.
- Apart from the model, validation of the data fed. Bias can enter through data collection and preparation before the model even runs.
- Deliberately testing the model against adversarial or edge case profiles to see where it breaks down or produces anomalous results.
- Benchmarking against existing traditional model approach.

Human oversight controls

Measures



Model Governance and Regulatory Accountability

- The Respective function owner should hold formal sign-off authority prior to model deployment and must be notified of any material changes.
- Customer grievance and redressal processes for review of disputed decisions.
- Maintain model documentation that supports regulatory examination, including data sources, variable selection rationale, validation results, and fairness testing outcomes.

Human oversight controls

Measures



Ongoing Monitoring and Continuous Improvement

- Regular monitoring for model drift, demographic changes, emerging biases due to changes in new business mix, etc.
- Investigation of significant deviations from expected outcomes.
- Periodic recalibration or retraining, triggered by pre-defined thresholds, such as loss ratio deviations or introduction of new parameters.
- Enhancement of the model over time to reduce manual referrals and increase straight-through processing.

Case Study - Governance for AI-Based Models

Agenda



The Board asks the Appointed Actuary to make a presentation covering the following aspects:

- How should fairness and bias be tested in AI-based underwriting and pricing models?
- What human oversight controls are necessary for AI-driven decisions?
- **What role should the Board play in AI governance and oversight and to establish AI Ethics?**
- Should the model be suspended pending remediation? What factors should be considered?
- What is the evolving role of the actuary in governance of AI-based decision systems?

Board's Governance Role

Context



- Traditional underwriting used a fixed rulebook where defined variables are mapped to pricing buckets; slow but fully transparent
- The AI engine draws on thousands of data points and behavioural signals simultaneously, finding patterns no manual process could replicate
- This leads to faster decisions, sharper risk selection, improved profitability
- AI has replaced our fixed rules with fluid, automated decisions
- Our role now is to build the governance infrastructure: dashboards, validation frameworks, oversight mechanisms in order to audit and control these automated decisions

Board's Governance Role

AI Dimension



What AI changes in the governance landscape

- Governance frameworks must incorporate fairness, explainability, model drift and proxy discrimination considerations
- Decision-making increasingly operates through automated and continuously evolving models
- Pricing and underwriting outcomes are generated at much greater scale and speed

The Board's role remains unchanged in principle

- Establish framework
- Set appetite
- Ensure accountability
- Demand information
- Exercise oversight

Board's Role in AI Governance

Committee and Policy



Committee and Policy

- Board to establish a dedicated AI Governance Committee and approve a formal AI Governance Policy
- Suggested members: Chief Risk Officer, Chief Technology Officer, Appointed Actuary, Head of Underwriting, Legal, Compliance, Finance and Internal Audit
- Committee will be responsible for oversight of model deployment, validation findings, governance monitoring, remediation tracking
- Policy establishes governance boundaries within which AI systems operate

Board's Role in AI Governance

Risk Appetite and Accountability

AI Risk Appetite

- Risk appetite framework should clearly define the boundaries within which automated underwriting can operate, including thresholds for high sum assured cases, exceptional medical conditions and defined uninsurable scenarios
- Cases exceeding defined thresholds, high-risk or exceptional cases should trigger mandatory manual review and approval

Accountability

- Every AI model should have a clearly identified accountability owner.
- Accountability should cover model performance, governance compliance, escalation and remediation
- Accountability mapping should be reviewed periodically and updated for model or organisational changes

Board's Role in AI Governance

Validation and Internal Reporting



Independent Model Validation or Peer review

- AI models should undergo independent validation or peer review
- Its scope should include model performance, data integrity, stability, fairness and governance controls
- Validation findings should be reviewed by committee and tracked to closure

AI Governance Internal Reporting

- Board and Governance Committee should receive structured reporting on model performance and governance indicators
- Reporting should present information in a business and governance context rather than only technical metrics
- Reporting should highlight areas requiring management attention, monitoring or escalation

Board's Role in AI Governance

Regulatory Preparedness and Remediation Oversight

Regulatory Preparedness

- AI governance framework should be in compliance with relevant IRDAI regulations and guidelines as issued, time to time
- Existing IRDAI requirements around fair and non-discriminatory pricing continue to apply to AI-driven systems
- Organisation should maintain readiness for evolving AI-related regulatory expectations and supervisory reviews

Remediation Oversight

- Board should oversee remediation plans for identified model or governance issues
- Remediation tracking should include accountable owners, timelines and status updates
- Material issues should be escalated through defined governance channels until closure

Board's Role in AI Oversight

Overview



The Board's oversight is exercised through the AI Governance Committee. The Board sets the expectations, receives Committee reports, and ensures that the oversight function is operating effectively.

The AI Governance Committee monitors the following in order to report to Board:

- Auto-Underwriting Volumes and Trends
- Actual vs Expected Experience
- Profitability
- Processing Speed
- Protocol Effectiveness

Board's Oversight in Practice

Day-to-day working and incidence reporting



Day-to-day working

- Review governance reporting with appropriate challenge and oversight
- Monitor emerging trends and key indicators
- Ensure consistency between management reporting and independent governance assessments
- Support continuous improvement in governance and operational effectiveness

Incidence reporting

- Receive complete and timely information on identified issues
- Independent review or auditor involvement should be initiated depending on materiality of the issue
- Priority focus on customer impact assessment, complaint resolution and corrective action; with the specific objective of resolving concerns before escalation to the Ombudsman
- Reputational risk should be managed actively

AI Ethics: Principles and Governance

Consequences



The Board defines and approves AI ethics principles that carry operational consequence

A. Fairness

AI outputs must not systematically disadvantage any customer group

Applies equally to AI-driven pricing

Fairness testing mandatory pre-deployment; continuous monitoring post-deployment; results reported to Board

B. Accountability

Every AI-driven decision should have an accountability owner

Explicit accountability map maintained; no ambiguity when a model causes harm

AI Ethics: Principles and Governance

Consequences



C. Transparency

Customers are entitled to a meaningful explanation

APS 34 Section 3.1 communication standard extends to how AI-driven outcomes are explained to policyholders

D. Proportionality

AI systems used only within validated scope; any extension requires revalidation and formal reapproval

APS 34 Section 2.10.3 - actuary must understand conditions under which a model is appropriate for use

E. Auditability

Every decision must be reconstructable: model version, inputs, output, and any override applied

There should be an adequate audit trail

Case Study - Governance for AI-Based Models

Agenda



The Board asks the Appointed Actuary to make a presentation covering the following aspects:

- How should fairness and bias be tested in AI-based underwriting and pricing models?
- What human oversight controls are necessary for AI-driven decisions?
- What role should the Board play in AI governance and oversight and to establish AI Ethics?
- **Should the model be suspended pending remediation? What factors should be considered?**
- What is the evolving role of the actuary in governance of AI-based decision systems?

Remediation

Trigger Factors & Remediation Decision Framework

Scope

Approach

A: Minimal & Targeted

- Targeted variables
- AI engines unchanged
- Pricing granularity & efficiency preserved

- Remove proxy variable or re-weight data.
- Corrected AI runs

B: Partial Suspension

- Suspend for biased cohorts
- Clean segments run uninterrupted
- Continuous Monitoring & drift alerts on

- Segregate AI and traditional workflows
- Audit trails mandatory
- Time-bound remediation plan

C: Full Suspension

- Entire model halted
- Business run on traditional models
- AI efficiency and pricing gains lost

- Full model suspended
- Written respective function owner opinion issued to Board
- Independent external validation required before re-deployment

Remediation

Trigger Factors & Remediation Decision Framework

Outcome

Impact on PH's

A: Minimal & Targeted

- Full pricing granularity retained.
- Operational efficiency preserved
- Post-correction model becomes robust

- No disruption
- All policies continue
- Human review overlay on affected cohorts as compensating control

B: Partial Suspension

- Partial efficiency loss
- Biased cohort via traditional UW/pricing models
- Escalation to Option C if remediation fails

- In-force protected
- Benefits and premiums unchanged
- New business for biased cohorts reverts to traditional UW

C: Full Suspension

- AI benefits nullified
- Redeployed model post correction will be fairer
- Post remediation restores pricing advantage

- PH's right to explanation of adverse AI decision
- Assess inflated-premium
- Excess premiums are refunded

Remediation

Trigger Factors & Remediation Decision Framework

A: Minimal & Targeted

- Stress-test reserves against bias scenarios
- Note sensitivity without qualification if below materiality

B: Partial Suspension

- Segmented reserve review: biased vs clean cohorts.
- Note sensitivity with qualification for biased segments

C: Full Suspension

- Full reserve re-estimation and strengthening of reserves
- Material impact disclosed to Regulator

Reserving

Solvency, Capital & Profitability

- Solvency is not affected
- Temporary capital buffer
- Profitability unaffected

- Solvency recalculation triggered if reserve > materiality threshold
- Capital buffer required
- Profitability affected - dual track costs & partial loss of AI economics

- Solvency may breach threshold
- Capital infusion necessary
- Profitability is severely affected due to suspension

Common Framework

Redeployment & Risk Register



01 Re-deployment Gate

▶ Actuary's Action

Resuming AI engines at any level requires independent external validation before going live.

▶ Outcome

Actuary documents validation scope, results and sign-off. Board approves re-deployment formally

02 Model Risk Register

▶ Actuary's Action

Decision rationale, test outcomes and escalation triggers recorded at each stage for all options

▶ Outcome

Audit trail: all decisions, model versions, bias test results and override justifications

Case Study - Governance for AI-Based Models

Agenda



The Board asks the Appointed Actuary to make a presentation covering the following aspects:

- How should fairness and bias be tested in AI-based underwriting and pricing models?
- What human oversight controls are necessary for AI-driven decisions?
- What role should the Board play in AI governance and oversight and to establish AI Ethics?
- Should the model be suspended pending remediation? What factors should be considered?
- **What is the evolving role of the actuary in governance of AI-based decision systems?**

Actuary's New Role

Fairness & enhanced role in Assumptions pricing & UW



01

Defining Fairness Before Training

► Actuary's Role

Defines fairness metric & demographic parity pre-build
Bias thresholds set to regulatory tolerance before model training begins

► Outcome

Fairness standard documented
Thresholds in model risk register

02

Enhanced UW/Pricing Governance

► Actuary's Role

Embedded in full UW & Pricing framework:
acceptance bands, sub-standard loading & override escalation

► Outcome

Enhanced role in UW and pricing governance

03

Assumption Alignment

► Actuary's Role

Certify AI segmentation consistent with mortality, morbidity and lapse assumptions

► Outcome

Improves fairness and reduces biases in pricing

Actuary's New Role

Granularity, persistency and mortality monitoring



04

Granular Experience Analysis

▶ Actuary's Role

High-frequency experience investigations by risk band using AI-generated records

▶ Outcome

Responsive assumption updates
Improves pricing and profitability

05

Persistency Management

▶ Actuary's Role

Identify lapse-prone policyholders with precision; validate persistency vs AI predictions

▶ Outcome

Reduced adverse selection risk
Adverse trends escalated to Board

06

Mortality Monitoring

▶ Actuary's Role

Near-real-time mortality and longevity monitoring enables earlier reserve adjustments

▶ Outcome

Continuous reserve adequacy
Threshold-triggered review and Regulatory compliance

Actuary's New Role

Return, compliance and Audit Trails



07

Periodic to Continuous Compliance

▶ Actuary's Role

Design real-time dashboards with AI drift thresholds that trigger automatic review

▶ Outcome

Proactive drift management
No regulatory surprises, outcomes documented and reported

08

Periodic Returns to Live Integration

▶ Actuary's Role

Near-term: enhanced periodic returns with AI metrics
Medium-term: real-time Regulatory integration

▶ Outcome

Seamless regulatory visibility into AI model performance
Live querying by IRDAI

09

Certifying AI Audit Trails

▶ Actuary's Role

Certify AI decision records (inputs, model version, output, confidence score, override).

▶ Outcome

Complete certified AI audit trail. Enforcement-ready- access provided to regulator on demand.

Actuary's New Role

Participate evolving Regulatory Standards & Solvency



10

Co-designing Regulatory Standards

► Actuary's Role

Contribute to IAI & IRDAI consultations to co-design AI governance & bias standards

► Outcome

Regulator-endorsed AI governance standards for Indian life insurance Industry

11

Evolving the Solvency Framework

► Actuary's Role

Contribute to evolving India's RBC framework to capture model risk & concentration risk

► Outcome

Updated solvency framework captures AI-driven risks, enhance protection of policyholder interests

Thank You | Q & A