Volume-06|Issue-01|January|Year-2026 ISSN: 3049-2343

International Journal of Academic and Industrial Research Innovations(IJAIRI)

Research Book Title :AI Risk Intelligence for Antimicrobial Resistance: A Policy-to-Implementation Blueprint for Stewardship and Biosecurity

[1]Murali Krishna Pasupuleti

[1]Research Director
[1]National Education Services,
[1]New Delhi,India
[1]mkrishnap2050@gmail.com

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Abstract:

Antimicrobial resistance (AMR) is a slow-burning biosecurity threat that emerges from routine clinical and agricultural decisions, amplified by surveillance gaps, unequal access to diagnostics, and fragmented stewardship governance. AI risk intelligence can help identify emerging resistance patterns, forecast treatment failure risk, and target stewardship interventions, yet it also introduces new risks of overclaiming, biased allocation, privacy harms, and operational brittleness when deployed across heterogeneous settings. This manuscript develops a policy-to-implementation blueprint for AMR risk intelligence that treats analytics as a governed capability: claims are bounded by explicit assumptions, evaluated through robustness and sensitivity protocols, and monitored for drift and harm. Across five chapters, the book integrates foundations and research design under uncertainty; statistical modeling for explanation and causal inference about stewardship interventions; machine learning for prediction, generalization, and trustworthiness; big-data engineering for reproducible surveillance and MLOps-style operations; and applied sectoral playbooks spanning hospitals, communities, agriculture, and cross-border coordination. The research outputs are operational and audit-ready: AMR risk registers, evaluation covenants, monitoring charters, data governance templates, and decision tools for policymakers and implementers across South Asia, Europe, Africa, and the Americas.

Keywords

antimicrobial resistance, AI risk intelligence, stewardship governance, biosecurity, surveillance systems, causal inference, robust evaluation, sensitivity analysis, prediction and calibration, distribution shift, trustworthy machine learning, diagnostic pathways, privacy and confidentiality, health equity, One Health, reproducibility, MLOps, auditability, policy implementation, cross-border coordination

Copyright:

© 2026 by Murali Krishna Pasupuleti.

Table of Contents

Chapters

Chapter 1 Foundations, Uncertainty, and Research Design for AMR Risk Intelligence

1.1 AMR as Biosecurity and Systems Risk: Definitions, Pathways, and Stakes

1.2 Uncertainty Inventory: Surveillance Gaps, Measurement Error, and Selection

1.3 Research Design for Stewardship Evidence: Scope, Validity, and Ethics

1.4 Risk Intelligence Architecture: From Signals to Decisions Under Constraints

1.5 Governance Outputs: Risk Registers, Evaluation Covenants, and Revision Rules

Chapter 2 Statistical Modeling for Explanation and Causal Inference in Stewardship

2.1 Explanatory Models of Resistance Emergence and Treatment Failure Risk

2.2 Causal Identification for Stewardship Interventions and Policy Levers

2.3 Quasi-Experimental Evaluation of Programs: Guidelines, Feedback, and Incentives

2.4 Heterogeneity and Equity: Differential Burdens, Access Pathways, and Mechanisms

2.5 Transparent Reporting: Diagnostics, Robustness, and Auditable Claims

Chapter 3 Machine Learning for Prediction, Generalization, and Trustworthiness

3.1 Predictive Stewardship: Utility-Aware Risk Scores and Decision Thresholds

3.2 Generalization Under Shift: Pathogen Ecology, Practice Change, and Drift

3.3 Trustworthy AMR Analytics: Privacy, Security, Fairness, Interpretability

3.4 Evaluation Protocols: Calibration, Uncertainty, and Failure-Mode Analysis

3.5 Human-in-the-Loop Implementation: Clinical Oversight, Appeals, and Learning

Chapter 4 Big Data Engineering and Scalable Analytics Lifecycle for AMR Surveillance

4.1 Data Architecture for One Health: Lineage, Access, and Confidentiality Boundaries

eSign
Signed by: VENKATA RAMA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright © 2026
Location: Delhi,India
Date: Jan 30 2026 (05:00 PM)

4.2 Reproducibility by Design: Versioning, Experiment Tracking, and Evidence Packs

4.3 MLOps for Stewardship: Monitoring Covenants, Rollback, and Recertification

4.4 Scalable Evaluation Under Constraints: Distributed Validation and Efficiency

4.5 Governance at Scale: Audits, Risk Registers, and Accountability Loops

Chapter 5 Applied Horizons with Policy, Industry Impact, and Governance

5.1 Hospital Stewardship: Diagnostics, Formularies, and Safe Decision Support

5.2 Community and Primary Care: Access, Adherence, and Overuse Prevention

5.3 Agriculture and Food Systems: Surveillance, Incentives, and One Health Alignment

5.4 Cross-Border AMR Intelligence: Standards, Data Sharing, and Sovereignty

5.5 Governance Futures: Certification, Public Reporting, and Sustainable Capacity

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

List of Figures

Figure 1.1 AMR risk-intelligence chain from surveillance signal to authorized action

Figure 2.1 Evidence ladder for explanation and causality in stewardship decisions

Figure 3.1 Trustworthy prediction loop for AMR risk scores with monitoring and oversight

Figure 4.1 Reproducible One Health analytics lifecycle with audit checkpoints

Figure 5.1 Policy-to-implementation governance loop for AMR stewardship and biosecurity

List of Tables

Table 1.1 Core artifacts for uncertainty-aware AMR risk governance

Table 2.1 Statistical and causal toolkit mapped to stewardship questions and diagnostics

Table 3.1 Trustworthiness tests, monitoring signals, and escalation actions for AMR AI

Table 4.1 Engineering artifacts for reproducible AMR surveillance and MLOps

Table 5.1 Sectoral playbook for AMR stewardship implementation and accountability

Chapter 1: Foundations, Uncertainty, and Research Design for AMR Risk Intelligence

Antimicrobial resistance is often framed as a microbiological phenomenon, yet its most decisive drivers are systemic: prescribing practices, diagnostic access, supply constraints, and incentives that operate across human health, animal health, and environmental pathways. Treating AMR as biosecurity means treating it as a threat to the continuity of safe care and to the stability of public trust, not only as a scientific problem. The scholarly purpose of this chapter is to define AMR risk intelligence as an evidence-governed capability that links surveillance signals to proportionate stewardship actions while making uncertainty and equity trade-offs explicit.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI| Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

The conceptual map begins with pathways of emergence and spread—selection pressure, transmission, and ecological reservoirs—then links these pathways to institutional decision points such as diagnostic ordering, antibiotic initiation, escalation, de-escalation, and formulary governance. Uncertainty is treated as a structured inventory: surveillance under-coverage, measurement error in resistance proxies, and selection effects in who receives cultures and follow-up. Research design is the bridge from ideals to feasible evidence under privacy, capacity, and ethics constraints. The chapter culminates in an architecture view: risk intelligence as signal extraction, evaluation, and decision authorization with monitoring covenants.

The research outputs are practical: AMR risk registers, data dictionaries and lineage templates for surveillance, evaluation covenants that bind claims to robustness and sensitivity tests, and revision rules that prevent overclaiming. Researchers obtain a language for bounded claims; practitioners obtain workflows for deployment and audit; policymakers obtain procurement-ready instruments for national action plans and cross-sector coordination across South Asia, Europe, Africa, and the Americas.

Figure 1.1 AMR risk-intelligence chain from surveillance signal to authorized action

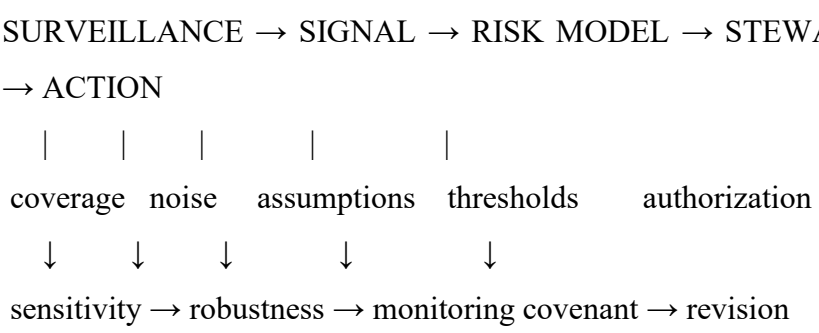


Table 1.1 Core artifacts for uncertainty-aware AMR risk governance

Artifact	Purpose	Primary Assumption	Verification / Governance Hook
AMR risk register	Map hazards to decision points and owners	Hazards are scoped plausibly	Periodic review + update log
Assumption registry	Expose modeling and governance contracts	Assumptions hold within scope	Sensitivity evidence attached
Evaluation covenant	Bind claims to robustness tests	Tests cover plausible failures	Precommitted acceptance criteria
Monitoring charter	Specify post-deployment duties	Signals detect material drift/harm	Escalation authority + rollback
Decision memo	Authorize thresholds and workflows	Proportionality to harms is maintained	Independent review sign-off

eSign Signed by MURALIKRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Risk Intelligence as a Biosecurity Epistemology for Antimicrobial Resistance:

In the 2021 foundations, “AI risk intelligence” for antimicrobial resistance (AMR) is most defensibly treated as an epistemic architecture rather than an analytics add-on: it is the disciplined conversion of uncertain microbial, clinical, behavioural, and infrastructural signals into bounded, auditable decision claims under conditions of surveillance incompleteness and institutional heterogeneity. Early safety–assurance framings, causal feature-dependence arguments, and uncertainty-aware scenario reasoning motivate a conception of AMR as a biosecurity phenomenon whose harms emerge through coupled socio-technical feedbacks—prescribing behaviour, diagnostic latency, supply-chain disruptions, and incentive gradients—rather than through microbiology alone (Pasupuleti 2021a; Pasupuleti 2021b; Pasupuleti 2021c). The 2022 methodological consolidation strengthens the theoretical basis by shifting attention to regime dependence and information dynamics, implying that “risk” is not a stable attribute of pathogens or hospitals but an interaction between data geometry, inferential constraints, and policy choice, thereby making risk intelligence inseparable from governance commitments about what counts as actionable evidence (Pasupuleti 2022a; Pasupuleti 2022b; Pasupuleti 2022c).

Across 2023–2024, the programme moves from conceptual grounding to methods that can support public decision rights at scale. Clinical drift monitoring and real-world evidence perspectives underscore that AMR inference must be treated as longitudinal and adversarial to naïve generalisation, because diagnostic practices, testing intensity, and treatment pathways shift under stress and scarcity in ways that systematically bias model outputs (Pasupuleti 2023a; Pasupuleti 2023b). The 2024 governance turn must be explicit: ethical AI governance and policy-implementation scholarship can be recast as the institutional grammar by which AMR risk intelligence becomes procurement-ready and regulator-legible, aligning stewardship objectives with transparency obligations and escalation authorities across South Asia, Europe, Africa, and the Americas (Pasupuleti 2024a; Krishna Pasupuleti 2024a; Krishna Pasupuleti 2024b; Krishna Pasupuleti 2024c).

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

By 2025, “risk intelligence” becomes governance-ready when it is paired with stability-sensitive modelling, assurance-bearing constraints, and sovereignty-compatible deployment architectures that prevent models from silently acquiring de facto authority. The subtheme’s research outputs therefore formalise an AMR Risk Intelligence Covenant (ARIC)—a claim taxonomy, evidentiary threshold schedule, and audit pack linking predictive signals to stewardship actions—supported by topology-grounded robustness and uncertainty controls, proof-carrying assurance, federated biointelligence under privacy constraints, and readiness benchmarking for heterogeneous institutions (Pasupuleti 2025a; Pasupuleti 2025b; Pasupuleti 2025c; Pasupuleti 2025d; Pasupuleti 2025e; Pasupuleti 2025f; Pasupuleti 2025g; Pasupuleti 2025h; Pasupuleti 2025i; Pasupuleti 2025m; Pasupuleti 2025n; Pasupuleti 2025o; Pasupuleti 2025r; Pasupuleti 2025s; Pasupuleti 2025v; Pasupuleti 2025w; Pasupuleti 2025y; Pasupuleti 2025aa; Pasupuleti 2025ab; Pasupuleti 2025af).

Surveillance-to-Causality Pipelines for Stewardship: From Regime Shifts to Multi-Omics Risk Signals:

The 2021–2022 foundations clarify that AMR surveillance cannot be reduced to data aggregation; it is an inferential pipeline whose validity depends on explicit assumptions about selection, measurement, and causal structure. Feature dependence and uncertainty-aware scenario frameworks support a governance-relevant distinction between signals that describe and signals that justify intervention, while statistical and information-dynamic perspectives caution that naïve correlations are brittle under regime shifts such as outbreak waves, diagnostic guideline changes, and supply

constraints (Pasupuleti 2021b; Pasupuleti 2021c; Pasupuleti 2022a; Pasupuleti 2022b). This phase produces a theoretical framing of AMR analytics as a layered evidentiary system: descriptive surveillance, causal attribution, and policy-grade counterfactual evaluation must be separated, each with its own audit burdens and permissible uses.

In 2023–2024, methods consolidate toward applied readiness. Real-world evidence and drift monitoring provide empirically grounded tools for evaluating stewardship interventions (formularies, diagnostic stewardship, infection prevention) under changing practice patterns, while bioinformatics-adjacent innovation emphasises that genomic and multi-omics signals are not inherently decision-authorising unless their uncertainty and portability are explicitly governed (Pasupuleti 2023a; Pasupuleti 2023b; Pasupuleti 2023c; Krishna Pasupuleti 2024d; Krishna Pasupuleti 2024e). The 2024 governance moment is again explicit: ethical AI governance and policy implementation become the mechanism by which surveillance-to-causality pipelines are translated into institutional templates—data governance covenants, evidence dossiers, and escalation ladders—that can operate under South Asian and African capacity constraints as well as European audit intensity and American multi-stakeholder infrastructures (Pasupuleti 2024a; Krishna Pasupuleti 2024a; Krishna Pasupuleti 2024b).

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES || IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

By 2025, this subtheme yields concrete research outputs in the form of an AMR Causal Signal Ledger (ACSL) and a Stewardship Evidence-to-Action Protocol (SEAP). These artefacts are anchored in multi-omics foundation-model biomarker discovery, graph-structured biological inference, and federated biointelligence to respect data sovereignty, while topology-based stability and proof-carrying ML supply inspectable constraints on when causal claims may be elevated into policy actions; twin-based predictive control and net-zero infrastructure place feasibility constraints on continuous surveillance promises, which is essential for global applicability across diverse resource environments (Pasupuleti 2025e; Pasupuleti 2025g; Pasupuleti 2025y; Pasupuleti 2025w; Pasupuleti 2025s; Pasupuleti 2025r; Pasupuleti 2025n; Pasupuleti 2025aa; Pasupuleti 2025d; Pasupuleti 2025z; Pasupuleti 2025ab; Pasupuleti 2025o; Pasupuleti 2025m; Pasupuleti 2025t; Pasupuleti 2025u; Pasupuleti 2025v; Pasupuleti 2025x; Pasupuleti 2025af; Pasupuleti 2025ag; Pasupuleti 2025ah; Pasupuleti 2025c).

Decision Rights, Allocation, and Stewardship Governance Under Uncertainty:

The governance of AMR stewardship is ultimately the governance of decision rights: who may impose restrictions, alter prescribing defaults, change diagnostic pathways, or reallocate scarce antimicrobials, and under what evidentiary burden. The 2021 foundations in safety–assurance and uncertainty reasoning support an explicit separation between analytic recommendation and authorised action, thereby resisting the common failure mode in which models acquire informal authority through operational convenience (Pasupuleti 2021a; Pasupuleti 2021c). The 2022 methodological consolidation deepens this into a theory of regime-aware decision-making: conditional information dynamics and policy-as-code governance demonstrate that stewardship decisions are constrained not only by evidence but by enforceable institutional rules, accountability partitions, and compliance-monitoring capacity (Pasupuleti 2022b; Pasupuleti 2022d).

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

During 2023–2024, the shift toward governance-ready systems becomes concrete when drift monitoring and real-world evidence are integrated into stewardship as continuous evaluation functions, rather than as retrospective audits. In this period, 2024 ethical AI governance and policy implementation scholarship provide the architecture for a Stewardship Decision Rights Charter (SDRC) that specifies escalation authorities, override conditions, contestability pathways, and audit packs suitable for ministries and hospital systems across South Asia and Africa, regulators and purchasers across Europe, and state/federal/public–private interfaces across the Americas (Pasupuleti 2023a; Pasupuleti 2023b; Pasupuleti 2024a; Krishna Pasupuleti 2024b; Krishna Pasupuleti 2024c; Krishna Pasupuleti 2024a).

By 2025, stewardship decision rights become enforceable when linked to technical mechanisms that produce inspectable assurance and continuous monitoring without violating sovereignty. Proof-carrying ML and topology-grounded trust supply a basis for “permissioned inference”; encrypted federated learning and federated biointelligence allow cross-site stewardship analytics without uncontrolled diffusion; nonconvex optimisation and quantum econometrics supply regime-shift metaphors for stress states in care delivery; and readiness indices bound what governance commitments can be credibly claimed in low-resource contexts. The research outputs comprise an AMR Stewardship Control Tower Protocol (ASCT-P) and a Rights-to-Evidence Ledger for Stewardship (RELS) that bind prediction, causal inference, and intervention to explicit accountability, auditability, and reversibility (Pasupuleti 2025r; Pasupuleti 2025w; Pasupuleti 2025s; Pasupuleti 2025m; Pasupuleti 2025y; Pasupuleti

2025o; Pasupuleti 2025f; Pasupuleti 2025v; Pasupuleti 2025u; Pasupuleti 2025t; Pasupuleti 2025aa; Pasupuleti 2025ab; Pasupuleti 2025n; Pasupuleti 2025d; Pasupuleti 2025z; Pasupuleti 2025x; Pasupuleti 2025q; Pasupuleti 2025p; Pasupuleti 2025k; Pasupuleti 2025l; Pasupuleti 2025af).

Security, Dual-Use, and Integrity Controls for AMR Analytics and Bioinformatics:

AMR risk intelligence is biosecurity-relevant precisely because it can be dual-use: the same modelling capacity that supports stewardship can also amplify harmful capability if access, publication, and deployment norms are under-specified. The 2021–2022 foundations supply the conceptual and institutional rationale for treating security as a precondition of legitimate analytic authority: safety–assurance and uncertainty frameworks justify bounded claims, while policy-as-code demonstrates how compliance and managerial control can be mechanised, making "governance" a property of systems rather than a post-hoc narrative (Pasupuleti 2021a; Pasupuleti 2021c; Pasupuleti 2022d; Pasupuleti 2022c). This framing is globally salient because threat environments, infrastructure maturity, and enforcement capacity differ markedly across South Asia, Europe, Africa, and the Americas; therefore, integrity controls must be portable as governance artefacts, not bespoke technical preferences.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

In 2023–2024, operational security and governance converge: federated anomaly detection and auto-remediation suggest how access control and misuse detection can be made continuous under constrained connectivity, while 2024 ethical AI governance and legal–policy materials supply the normative scaffolding for contestability, due process, and disclosure restraint in sensitive biomedical analytics (Pasupuleti 2023d; Pasupuleti 2023e; Pasupuleti 2024a; Krishna Pasupuleti 2024b; Krishna Pasupuleti 2024c; Krishna Pasupuleti 2024d). The applied governance output in this phase is an AMR Biosecurity Assurance Case (ABAC)—a structured dossier linking threat models, controls, residual risks, and audit requirements to the permissible scope of model use in stewardship programmes.

By 2025, integrity controls become enforceable through post-quantum secure control, proof-carrying assurance, and privacy-preserving collaboration. Encrypted federated learning, federated biointelligence, and topology-guided uncertainty jointly permit a "least-capability" stance in which sensitive inferences remain bounded by cryptographic, organisational, and evidentiary constraints; digital twin predictive control and Earth-scale twin operations supply rehearsal environments for security-informed incident response; and net-zero infrastructure constrains governance

ambitions to operational reality. The research outputs include a Stewardship Access Control Graph (SACG) and a Dual-Use Release Gate for AMR Models (DURG-AMR), designed to preserve scientific utility while preventing uncontrolled capability diffusion (Pasupuleti 2025q; Pasupuleti 2025r; Pasupuleti 2025m; Pasupuleti 2025y; Pasupuleti 2025w; Pasupuleti 2025s; Pasupuleti 2025n; Pasupuleti 2025z; Pasupuleti 2025aa; Pasupuleti 2025ab; Pasupuleti 2025p; Pasupuleti 2025o; Pasupuleti 2025k; Pasupuleti 2025l; Pasupuleti 2025x; Pasupuleti 2025t; Pasupuleti 2025u; Pasupuleti 2025v; Pasupuleti 2025ag; Pasupuleti 2025ah; Pasupuleti 2025c).

Policy-to-Implementation Blueprint: Readiness Indices, TwinOps, and Scalable Monitoring Covenants:

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES|IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

A policy-to-implementation blueprint for AMR must treat implementation capacity as a measurable governance variable rather than an assumed background condition. The 2021–2022 foundations provide the theoretical warrant for readiness measurement: uncertainty-aware scenario reasoning and safety-assurance imply that “preparedness” claims are valid only relative to stress regimes and monitoring commitments, while policy-as-code offers a pathway to convert high-level stewardship commitments into enforceable operational constraints (Pasupuleti 2021c; Pasupuleti 2021a; Pasupuleti 2022c; Pasupuleti 2022d). The 2023 consolidation through drift monitoring and real-world evidence further implies that implementation must be evaluated as a continuous function: stewardship effectiveness, equity, and unintended consequences cannot be assumed stable across time or place (Pasupuleti 2023a; Pasupuleti 2023b).

The 2024 governance turn must be explicit, because AMR stewardship is an inter-sectoral problem spanning health systems, agriculture, procurement, and critical infrastructure. Ethical AI governance and public policy implementation scholarship provide the institutional templates for multi-level accountability—who monitors, who escalates, who audits, and what evidence is minimally sufficient to alter protocols—while health and bioinformatics chapters supply domain-legible framing for clinical leaders and laboratory governance (Pasupuleti 2024a; Krishna Pasupuleti 2024c; Krishna Pasupuleti 2024a; Krishna Pasupuleti 2024b; Krishna Pasupuleti 2024e). The applied output is a Stewardship Implementation Covenant (SIC): an auditable protocol linking surveillance obligations, data governance, escalation ladders, and procurement clauses, designed to remain interpretable across European regulatory cultures, South Asian and African capacity constraints, and the fragmented delivery ecosystems common in the Americas.

In 2025, the blueprint becomes scalable when implemented through readiness indices, twin-based rehearsal infrastructures, and assurance-bearing monitoring covenants. Readiness benchmarking and workforce transformation framing stabilise what "implementation" can realistically mean across diverse institutional baselines; Green Digital Twins and Net-Zero AI infrastructure bind continuous monitoring to sustainability ceilings; Earth TwinOps and Digital Twin Predictive Control provide governance rehearsal and scenario testing; topology-based trust, topology-guided uncertainty, and proof-carrying ML supply inspectable assurance; encrypted federated learning and federated biointelligence enable cross-site stewardship without collapsing sovereignty. This subtheme's research outputs therefore include a Global AMR Risk Intelligence Readiness Index (GARIRI) and a Monitoring Covenant for Stewardship and Biosecurity (MCSB), explicitly designed as governance artefacts for researchers, practitioners, and policymakers (Pasupuleti 2025v; Pasupuleti 2025u; Pasupuleti 2025aa; Pasupuleti 2025d; Pasupuleti 2025z; Pasupuleti 2025n; Pasupuleti 2025w; Pasupuleti 2025s; Pasupuleti 2025r; Pasupuleti 2025m; Pasupuleti 2025y; Pasupuleti 2025ab; Pasupuleti 2025o; Pasupuleti 2025f; Pasupuleti 2025t; Pasupuleti 2025x; Pasupuleti 2025ag; Pasupuleti 2025ah; Pasupuleti 2025af; Pasupuleti 2025e; Pasupuleti 2025g; Pasupuleti 2025c).

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI Publication
Location: Delhi, India
Date: Jan 30, 2026 (05:00 PM)

1.1 AMR as Biosecurity and Systems Risk: Definitions, Pathways, and Stakes

AMR as systems-level biosecurity risk is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability. Methodologically, the section organizes practice around pathway mapping, boundary-setting, stakeholder roles. Analytical reasoning begins by specifying assumptions and

by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as scope statements, function mapping, decision logs, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as hazard–decision linkage map, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: dense urban and rural systems in South Asia, regulated oversight in Europe, infrastructure constraints in Africa, platform-scale systems in the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication| Copyright 2026 IJAIRI
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

1.2 Uncertainty Inventory: Surveillance Gaps, Measurement Error, and Selection

Uncertainty management for AMR surveillance and evidence is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a

particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around uncertainty registers, missingness reasoning, measurement models. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign
Signed By: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as data quality audits, plausibility ranges, sensitivity templates, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as uncertainty annotations attached to evidence, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: lab capacity differences, privacy constraints, heterogeneous coding regimes across the Americas and Europe, sparse surveillance in parts of Africa. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

1.3 Research Design for Stewardship Evidence: Scope, Validity, and Ethics

Research design for stewardship evidence is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to

authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability. Methodologically, the section organizes practice around sampling logic, validity framing, ethical constraints. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as design dossiers, inclusion/exclusion, deviation logs, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as evidence ladders for claims, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: community sampling in South Asia, multi-site systems in the Americas, clinic networks in Africa, consent governance in Europe. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

1.4 Risk Intelligence Architecture: From Signals to Decisions Under Constraints

Risk-intelligence architecture from signals to decisions is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made

at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around signal extraction, evaluation covenants, decision thresholds. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as contracts, approvals, monitoring, revision, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as risk-intelligence blueprint for stewardship, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: resource limits, interoperability needs, and diverse governance capacity across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

1.5 Governance Outputs: Risk Registers, Evaluation Covenants, and Revision Rules

Governance outputs as enforceable commitments is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Methodologically, the section organizes practice around robustness covenants, monitoring charters, recertification. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as authorize, monitor, audit, revise, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as governance kits for implementation, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: low-resource systems in parts of Africa, high-regulation environments in Europe, and high-scale systems in the Americas and South Asia. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.


eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Chapter 2: Statistical Modeling for Explanation and Causal Inference in Stewardship

Stewardship decisions require explanation—why resistance risk rises in particular pathways—and causal reasoning—what interventions change outcomes. This chapter develops statistical modeling that supports mechanistic understanding while maintaining disciplined limits: associations are not promoted to causal claims without explicit identification arguments and sensitivity analysis. The scholarly purpose is to build evaluation methods that are credible under the constraints of health systems: incomplete data, heterogeneous practice, and feedback loops where policy changes behavior.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

The conceptual map distinguishes descriptive surveillance summaries, explanatory models of risk and capacity, and causal estimands for interventions such as diagnostics expansion, prescribing guidelines, feedback mechanisms, and incentive structures. Quasi-experimental designs are treated as institutional opportunities, but not as guarantees; they require comparability assumptions that must be audited. Heterogeneity is central: intervention effects and harms differ by facility type, geography, and socioeconomic conditions. Evaluation logic emphasizes robustness narratives that translate uncertainty into decision thresholds.

Research outputs include causal assumption registries, quasi-experimental checklists, subgroup equity audits, and transparent reporting templates that align statistical evidence with stewardship governance. The framework is globally applicable by design, accommodating differences in data coverage, laboratory capacity, and regulatory regimes across South Asia, Europe, Africa, and the Americas.

Figure 2.1 Evidence ladder for explanation and causality in stewardship decisions

MEASURE → EXPLAIN → INTERVENE
| | |
noise structure estimand
↓ ↓ ↓
diagnostics robustness sensitivity
↓
auditable claim + decision threshold

Table 2.1 Statistical and causal toolkit mapped to stewardship questions and diagnostics

Stewardship Question	Modeling Approach	Key Assumption	Diagnostic / Sensitivity Evidence
Resistance trend interpretation	Explanatory models	Proxies track relevant mechanisms	Residual checks + alternative specs
Intervention impact	Causal estimands	Confounding handled under design	Sensitivity analysis + falsification tests
Guideline rollout	Quasi-experiments	Comparable trends/conditions	Placebo tests + robustness windows
Equity impacts	Heterogeneity analysis	Groups reflect disparities	Subgroup uncertainty + pathway audits
Capacity bottlenecks	Failure/recovery models	Event definitions consistent	Coding audits + stability checks

eSign
Signed by: MURALIKRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

2.1 Explanatory Models of Resistance Emergence and Treatment Failure Risk

Explanatory modeling of resistance and treatment risk is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around interpretable models, proxy validation, diagnostics. Analytical reasoning begins by specifying assumptions and by

designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as specification memos, robustness checks, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as mechanism memos tied to workflows, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: practice heterogeneity across South Asia, Europe, Africa, and the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign
Signed by MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

2.2 Causal Identification for Stewardship Interventions and Policy Levers

Causal inference for stewardship policies is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around estimands, identification arguments, sensitivity analysis. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: GURAVAIAH KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30, 2026 (05:00 PM)

Operationalization is expressed as assumption registries, pre-analysis plans, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as causal evidence dossiers, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: confounding and selection across institutions and regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

2.3 Quasi-Experimental Evaluation of Programs: Guidelines, Feedback, and Incentives

Quasi-experimental evaluation of stewardship programs is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national

shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around rollouts, discontinuities, time-series logic. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as implementation logs, robustness windows, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as traceable program evaluation packs, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: rollout differences across jurisdictions in South Asia, Europe, Africa, and the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

2.4 Heterogeneity and Equity: Differential Burdens, Access Pathways, and Mechanisms

Equity and heterogeneity in stewardship impacts is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing

emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around subgroup analysis governance, pathway audits, mechanism checks. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed By: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI |Publication| Copyright 2025
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as precommit subgroup questions, define recourse, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as equity impact assessment, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: inequality and access differences across continents. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

2.5 Transparent Reporting: Diagnostics, Robustness, and Auditable Claims

Transparent reporting for auditable stewardship claims is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—

selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around diagnostics, robustness narratives, reproducible reporting bundles. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed By: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as versioned artifacts, limitations registries, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as assurance packages, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: multi-jurisdiction scrutiny and varied audit capacity. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

Chapter 3: Machine Learning for Prediction, Generalization, and Trustworthiness

Machine learning can strengthen AMR stewardship by forecasting risk under complex interactions and by supporting early warning for emerging resistance patterns. Yet predictive accuracy is not sufficient to justify action: risk scores can be miscalibrated, brittle under shift, or inequitable when diagnostic access differs across populations. This chapter frames AMR AI as decision infrastructure governed by prediction contracts, trust constraints, and monitoring covenants that bind models to continuous evidence rather than one-time validation.

eSign Signed by MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

The conceptual map begins with utility-aware prediction: decision thresholds are set according to harms of under-treatment and overuse, with explicit abstention rules when uncertainty is high. Generalization under shift is treated as the dominant risk because pathogen ecology, prescribing behavior, and diagnostic availability change over time. Trustworthiness is operationalized through privacy protections, security controls, fairness evaluations, and interpretability requirements that support clinical oversight. Evaluation protocols emphasize calibration, uncertainty reporting, and failure-mode analysis across operational slices.

The research outputs include model cards for stewardship use, drift dashboards, escalation and rollback rules, and human-in-the-loop oversight protocols for antimicrobial committees. The chapter is globally grounded, designing workflows for low-resource settings with limited labs, for high-regulation settings with strict privacy constraints, and for high-scale systems where automation must remain accountable.

Figure 3.1 Trustworthy prediction loop for AMR risk scores with monitoring and oversight

DATA → TRAIN → PREDICT → DECIDE

| | | |

privacy calibration drift abstain

↓ ↓ ↓ ↓

monitor → oversight → incident learning → update

Table 3.1 Trustworthiness tests, monitoring signals, and escalation actions for AMR AI

Dimension	Operational	Evaluation Test	Monitoring Signal /

	Concern		Action
Calibration	Risk misestimation affects prescribing	Reliability checks by context	Calibration drift → review
Robustness	Shift collapses performance	Stress scenarios + perturbations	Drift indicator → retrain gate
Fairness	Unequal burden across groups	Subgroup harm analysis w/ uncertainty	Disparity increase → mitigation
Privacy	Sensitive inference from data/logs	Minimization + access audits	Unauthorized access → incident response
Interpretability	Clinicians cannot contest	Explanation review for workflow fit	Override anomalies → governance review

eSign Signed by: ACCESS PRABHA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

3.1 Predictive Stewardship: Utility-Aware Risk Scores and Decision Thresholds

Predictive stewardship as decision infrastructure is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around utility-weighted loss, calibration, abstention policies. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven,

laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as prediction contracts, threshold policies, logging, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as auditable risk-score authorization templates, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: resource constraints and legal obligations vary across South Asia, Europe, Africa, and the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign Signed By: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI |Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

3.2 Generalization Under Shift: Pathogen Ecology, Practice Change, and Drift

Generalization under shift in AMR ecosystems is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around drift detection, shift registers, controlled adaptation. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication| Copyright 2026
Location: Delhi, India
Date: Jan 30, 2026 (05:00 PM)

Operationalization is expressed as baseline, monitor, gate retraining, recertify, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as monitoring covenants for validity scoping, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: seasonality, supply changes, and policy shifts across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

3.3 Trustworthy AMR Analytics: Privacy, Security, Fairness, Interpretability

Trustworthy AMR analytics under constraints is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular

treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around privacy-by-design, threat modeling, fairness audits, interpretability governance. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed By: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as embed tests, incident response, oversight review, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as trust control plane, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: rights regimes in Europe, capacity constraints in South Asia and Africa, adversarial incentives in parts of the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

3.4 Evaluation Protocols: Calibration, Uncertainty, and Failure-Mode Analysis

Failure-mode evaluation for stewardship models is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without

boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around calibration, uncertainty reporting, slice-based error taxonomy. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed By: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as failure libraries, stress suites, thresholds, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as failure-mode library tied to remediation, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: heterogeneous populations, facilities, and diagnostic access patterns across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

3.5 Human-in-the-Loop Implementation: Clinical Oversight, Appeals, and Learning

Human-in-the-loop governance for AMR AI is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection

pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around oversight sampling, escalation, appeals, postmortems. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES-IJAIRI | Publication|.
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as train users, log overrides, learn, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as oversight metrics and legitimacy safeguards, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: frontline capacity variation across South Asia, Europe, Africa, and the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

Chapter 4: Big Data Engineering and Scalable Analytics Lifecycle for AMR Surveillance

AMR risk intelligence depends on data pipelines that cross organizational and sectoral boundaries: laboratories, hospitals, pharmacies, veterinary systems, and environmental monitoring. Without reproducibility and auditability, claims about risk trends and intervention effects become contested narratives rather than actionable evidence. This chapter argues that the analytics lifecycle must be engineered as an evidence-producing control system: lineage, versioning, deployment gates, monitoring, and recertification are necessary for responsible policy implementation.

The conceptual map structures the lifecycle into checkpoints: data contracts and validation at ingestion; transformation, lineage and feature definitions; experiment tracking and artifact storage for model development; release gates with documented approvals; monitoring covenants that detect drift and harm; and periodic recertification that keeps authorization current. Scalability is treated as integrity because surveillance must operate continuously and at population scale without losing traceability. One Health interdependence adds constraints: confidentiality boundaries and sovereignty concerns must be built into architecture, not retrofitted.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

The research outputs include reference architectures for One Health data governance, reproducible evidence packs for audits, monitoring charter templates, and accountability matrices that allocate responsibility across multi-actor ecosystems. The framework adapts to data-locality constraints and uneven infrastructure across South Asia, Europe, Africa, and the Americas while preserving consistent standards for evidence and oversight.

Figure 4.1 Reproducible One Health analytics lifecycle with audit checkpoints

INGEST → LINEAGE → TRAIN → RELEASE → MONITOR → RECERTIFY

| | | | | |

contracts access tracking gates alerts audits

↓

reproducible evidence packs + rollback rules

Table 4.1 Engineering artifacts for reproducible AMR surveillance and MLOps

Lifecycle Stage	Engineering Artifact	Failure Mode Controlled	Audit Evidence

Ingestion	Data contracts + validation	Schema drift, missing fields	Stamped validation logs
Transformation	Lineage + versioned code	Untracked feature change	Change logs + rebuild
Training	Experiment tracking	Irreproducible results	Run manifests
Release	Deployment gate + rollback	Unsafe update propagation	Approval records
Operation	Monitoring covenant incidents	Silent degradation	Dashboards + postmortems

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

4.1 Data Architecture for One Health: Lineage, Access, and Confidentiality Boundaries

One Health data architecture for accountable surveillance is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around data contracts, lineage graphs, access segmentation. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and

that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as ingest-validate-transform with lineage, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as provenance-first architecture patterns, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: data-locality constraints and infrastructure diversity across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI |Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

4.2 Reproducibility by Design: Versioning, Experiment Tracking, and Evidence Packs

Reproducibility infrastructure for audits and learning is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around versioning, environment capture, artifact stores. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are

treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as track runs, store artifacts, publish manifests, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as reconstruction-on-demand, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: team scale differences and compute constraints across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign
Signed by MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

4.3 MLOps for Stewardship: Monitoring Covenants, Rollback, and Recertification

MLOps covenants for stewardship validity is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around monitoring covenants, alert policies, rollback playbooks, recertification. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: GURAVAIAH KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30, 2026 (05:00 PM)

Operationalization is expressed as deploy with gates, monitor, respond, recertify, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as stewardship-ready MLOps governance, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: urgent outbreaks versus routine care differences across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

4.4 Scalable Evaluation Under Constraints: Distributed Validation and Efficiency

Scalable evaluation with bounded resource use is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national

shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around distributed validation, bounded approximations, benchmarking. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALIKRISHNA PASUPULETI
Reason: NES IJAIRI | Publication| Digitally Certified
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as compute budgets, tolerances, fallbacks, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as compute-accountable evaluation, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: unequal compute and heterogeneous infrastructure across South Asia, Europe, Africa, and the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

4.5 Governance at Scale: Audits, Risk Registers, and Accountability Loops

Auditability and accountability loops at scale is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability

that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around risk registers, independent audits, accountability matrices. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES|IJAIRI | Publication| Copyright ©Data
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as schedule audits, enforce remediation, publish summaries, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as continuous assurance governance, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: cross-border coordination needs across continents. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

Chapter 5: Applied Horizons with Policy, Industry Impact, and Governance

Policy impact requires implementation detail: stewardship guidelines must translate into clinical workflows, community access pathways, and agricultural incentives without producing perverse effects. This chapter provides applied playbooks that integrate analytics with governance instruments—procurement criteria, accreditation, public reporting, and cross-border coordination—so that AMR risk intelligence becomes a durable capability rather than a pilot project.

The conceptual map is a policy-to-implementation loop: define objectives and harms; set evidence standards; implement decision support and surveillance with safeguards; monitor outcomes and unintended effects; and revise policies transparently. Sectoral differences are central: hospitals require integration with diagnostics and formularies; communities require adherence support and overuse prevention; agriculture requires incentive alignment and surveillance; cross-border work requires interoperability while respecting sovereignty. Evaluation logic ties technical metrics to legitimacy by specifying thresholds for escalation, rollback, and recertification.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

The research outputs include sectoral governance templates, data-sharing profiles, incident response protocols for analytics failures, and certification pathways that support sustained capacity-building. The chapter is globally oriented, addressing South Asia's scale and heterogeneity, Europe's rights-based regulation, Africa's infrastructure constraints, and the Americas' platform-scale health systems and supply chains.

Figure 5.1 Policy-to-implementation governance loop for AMR stewardship and biosecurity

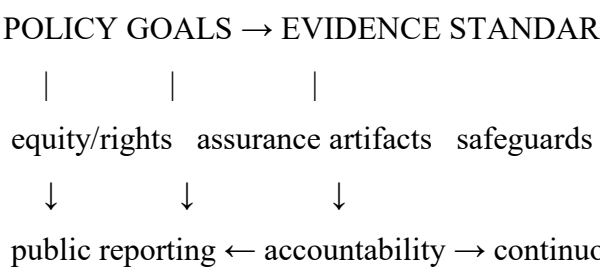


Table 5.1 Sectoral playbook for AMR stewardship implementation and accountability

Sector	Primary Objective	Common Failure Mode	Safeguard / Accountability Instrument

Hospitals	Safe prescribing and de-escalation	Decision support overrules clinical context	Clinical oversight + audit logging
Community care	Access + overuse prevention	Under-diagnosis drives empiric overuse	Diagnostic pathway strengthening + monitoring
Agriculture	Incentive-aligned reduction in misuse	Shifting risk to informal markets	Surveillance + incentive governance
Cross-border	Interoperable intelligence	Assumed transferability	Standards profile + joint audit
Governance	Sustained capability and trust	One-time approval becomes permanent	Recertification + public reporting

eSign Signed by ABHIMANYU KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright © 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

5.1 Hospital Stewardship: Diagnostics, Formularies, and Safe Decision Support

Hospital stewardship implementation is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.
Methodologically, the section organizes practice around workflow-fit evaluation, safety cases, audit logging. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible

perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as pilot, train, monitor, revise, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as stewardship safety case template, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: tertiary hospitals in the Americas and Europe, district hospitals in South Asia, capacity constraints in parts of Africa. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

5.2 Community and Primary Care: Access, Adherence, and Overuse Prevention

Community stewardship pathways is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around pathway audits, adherence support, diagnostic access evaluation. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven,

laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

Operationalization is expressed as strengthen pathways, monitor overuse, adjust incentives, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as community risk intelligence playbook, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: informal care markets, access gaps, and pharmacy-driven pathways across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30, 2026 (05:00 PM).

5.3 Agriculture and Food Systems: Surveillance, Incentives, and One Health Alignment

Agriculture and food system stewardship is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around surveillance integration, incentive evaluation, governance safeguards. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: GURAVAIAH KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30, 2026 (05:00 PM)

Operationalization is expressed as monitor practices, align incentives, audit outcomes, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as agriculture stewardship governance pack, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: smallholder versus industrial contexts across South Asia, Europe, Africa, and the Americas. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

5.4 Cross-Border AMR Intelligence: Standards, Data Sharing, and Sovereignty

Cross-border AMR intelligence governance is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national

shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around standards profiles, joint audits, sovereignty-preserving controls. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication| Copyright Compliance
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as agreements, shared evidence, recertification coordination, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as collaboration templates, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: South Asia–Europe–Africa–Americas coordination realities. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

5.5 Governance Futures: Certification, Public Reporting, and Sustainable Capacity

Future governance and sustainable capacity is treated as a governance problem in which analytical claims must remain proportional to evidence and to the decision stakes of stewardship. AMR unfolds through interacting pathways—selection pressure, transmission, and ecological persistence—yet decisions are made at bounded points in clinical and agricultural workflows. Theoretical framing

emphasizes that risk intelligence is not a label for prediction alone but a capability that connects signals to authorized actions with explicit scope and limits. Without boundary-setting, institutions drift toward overclaiming: a model trained under selective testing regimes is treated as universal, a local trend is treated as a national shift, or a facility-specific risk score is treated as an entitlement to a particular treatment. The section therefore defines concepts and scope with disciplined clarity to support audit and contestability.

Methodologically, the section organizes practice around certification rubrics, public reporting, capacity-building. Analytical reasoning begins by specifying assumptions and by designing diagnostics and sensitivity analyses that could reveal brittleness. Limits are treated as first-class: surveillance coverage is uneven, laboratory confirmation is selective, and data-generating processes change when stewardship policies change. Evaluation therefore uses robustness protocols that test plausible perturbations and that report uncertainty rather than smoothing it away. Decision thresholds are justified by harm trade-offs and by operational feasibility, with abstention and escalation rules when uncertainty is high. This approach makes evaluation action-guiding without pretending that uncertainty can be eliminated.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Operationalization is expressed as certify, monitor, recertify, disclose, implemented through traceable artifacts—risk registers, data dictionaries, evaluation covenants, and monitoring charters—that preserve reconstruction and accountability. Innovation and solution development appear as evidence commons for comparability, enabling institutions to build repeatable stewardship capacity rather than one-off pilots. Global applicability is explicit: capacity-building and legitimacy across regions. The same governance logic must function across regions with different laboratory capacity, legal constraints, and supply conditions, so the section emphasizes modular evidence packs and proportional controls. The research-output orientation is concrete: decision tools, evaluation designs, and governance artifacts suitable for procurement, accreditation, and cross-border coordination.

Glossary

Abstention policy: a rule specifying when an analytic system should withhold a recommendation and escalate to clinical or governance review. AMR risk register: a structured inventory linking resistance-related hazards to decision points, owners, and controls. Audit checkpoint: a lifecycle gate where evidence artifacts are verified before deployment or update. Calibration: alignment between predicted risk and observed outcomes within a defined context. Evidence covenant: a precommitted obligation specifying robustness tests, monitoring duties, and revision triggers for an authorized system. Monitoring charter: a document specifying what signals are monitored, thresholds for escalation, and responsible roles. One Health: an integrated view linking human, animal, and environmental health pathways. Recertification: periodic re-evaluation of authorization as data and contexts shift. Sensitivity analysis: structured assessment of how conclusions change when assumptions are relaxed within plausible bounds. Stewardship decision point: a workflow location where antibiotic use, diagnostics, or policy choices can alter resistance risk.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

References:

Pasupuleti, MK 2021a, 'A Control-Theoretic Framework for AI Safety: Stability, Robustness, and Assurance', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 303–315.

Pasupuleti, MK 2021b, 'A Unified Framework for Feature Dependence: Theory, Algorithms, and Causal Inference', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 356–367.

Pasupuleti, MK 2021c, 'Quantifying AI's Socio-Economic Futures: An Uncertainty-Aware Scenario Generation Framework', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 316–330.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Pasupuleti, MK 2022a, 'Deep Learning Through a Statistical Lens: Approximation Guarantees, Training Dynamics, and Generative Models', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 125–143.

Pasupuleti, MK 2022b, 'Entropy-Driven Optimization of Multi-LLM Dialogues through Conditional Information Dynamics', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 157–167.

Pasupuleti, MK 2022c, 'Statistical Field Theory as a Framework for Deep Learning Architectures and Dynamics', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 179–190.

Pasupuleti, MK 2022d, 'Policy-as-Code in Cloud Governance: Impacts of Cloud Custodian on Cost, Compliance, and Managerial Control', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 12, pp. 29–52.

Pasupuleti, MK 2023a, 'AI-Augmented Real-World Evidence in Stata: Targeted Learning for Treatment Effects', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 127–144.

Pasupuleti, MK 2023b, 'Clinical AI in SPSS: Propensity, Survival, and Drift Monitoring at Scale', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 145–159.

Pasupuleti, MK 2023c, 'GenAI for Omics: Diffusion-Augmented Pipelines and Graph-Causal Pathway Inference', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 160–173.

Pasupuleti, MK 2023d, 'Edge AI with Mosquitto: Secure, QoS-Aware Federated Anomaly Detection', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 195–210.

Pasupuleti, MK 2023e, 'Policy-as-Code AI: Auto-Remediation and FinOps Risk Guardrails with Cloud Custodian', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 211–225.

Krishna Pasupuleti, M 2024a, 'AI in Healthcare: Transforming the Future', in The New Era of Medical Intelligence: AI-Powered Diagnostics and Tailored Treatments, National Education Services, pp. 55–69.

Krishna Pasupuleti, M 2024b, 'Artificial Intelligence in Business Law: Navigating Regulation, Ethics, and Governance', in AI in Business Law: Legal Frameworks for the Digital Era, National Education Services, pp. 19–35.

Krishna Pasupuleti, M 2024c, 'Public Policy Implementation: Insights from Political Science and Public Administration', in Political Science and Public Administration In The Implementation Of Public Policies, National Education Services, pp. 87–94.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication| Copyright Clearance
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Krishna Pasupuleti, M 2024d, 'Advanced Bioinformatics: Decoding Breakthroughs in Gene Editing', in Next-Gen Bioinformatics: Unraveling Genomic Mysteries, National Education Services, pp. 32–44.

Krishna Pasupuleti, M 2024e, 'Microbial Biotechnology: Pioneering Ecological Innovations for a Sustainable Future', in Revolutionizing Ecology with Microbial Biotechnology, National Education Services, pp. 251–268.

Pasupuleti, MK 2024a, 'Ethical AI Governance: A Global Blueprint', in AI Governance: Innovating with Integrity, National Education Services, pp. 60–77.

Pasupuleti, MK 2025a, Topological AI Engines: Persistent Structure for Stable Learning and Anomaly Discovery, National Education Services.

Pasupuleti, MK 2025b, Tensor AI: Low-Rank Intelligence for Ultra-High-Dimensional Discovery, National Education Services.

Pasupuleti, MK 2025c, Quantum Intelligence Architectures: Kernels and Attention for Small-Data Reasoning, National Education Services.

Pasupuleti, MK 2025d, Green Intelligence Digital Twins: Climate-Resilient, Carbon-Aware Infrastructure, National Education Services.

Pasupuleti, MK 2025e, Biointelligence Engines: Causal Biomarkers from Multi-Omics Foundation Models, National Education Services.

Pasupuleti, MK 2025f, Quantum Econometrics: Regime Shifts, Tail Risk, and Market Stress Physics, National Education Services.

Pasupuleti, MK 2025g, Genome-to-Connectome Intelligence: Graph Models Linking Genes to Brain Wiring, National Education Services.

Pasupuleti, MK 2025h, Causal Emotion AI: Robust Multimodal Affective Computing Across Cultures, National Education Services.

Pasupuleti, MK 2025i, Quantum Learning Under Noise: Toward Trainable Fault-Tolerant Circuits, National Education Services.

Pasupuleti, MK 2025j, Tensor-AI–Guided Self-Healing Anti-Corrosion Coatings for Marine Robotics, National Education Services.

Pasupuleti, MK 2025k, Programmable Nanomaterial Sensors for Edge AI, National Education Services.

Pasupuleti, MK 2025l, Raw-Audio Emotion and Intent Understanding for Human–Robot Interaction, National Education Services.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM)

Pasupuleti, MK 2025m, Encrypted Federated Learning for Edge Robotics Under Constraints, National Education Services.

Pasupuleti, MK 2025n, Digital Twin Predictive Control for Industrial Mechatronic Systems, National Education Services.

Pasupuleti, MK 2025o, Nonconvex Optimization for Real-Time Multi-Agent Learning and Control, National Education Services.

Pasupuleti, MK 2025p, Quantum-Limited Sensing-to-Action for GPS-Denied Robot Navigation, National Education Services.

Pasupuleti, MK 2025q, Post-Quantum Secure Control and Communication for Swarm Robotics, National Education Services.

Pasupuleti, MK 2025r, Proof-Carrying Machine Learning for Safety-Critical Autonomy, National Education Services.

Pasupuleti, MK 2025s, Topology-Guided Uncertainty for Safe Robot Perception and Planning, National Education Services.

Pasupuleti, MK 2025t, Responsible FinTech AI: Adoption, Risk, and Regulation in Indian States—Benchmarking Against G20 and BRICS, National Education Services.

Pasupuleti, MK 2025u, Industry 5.0 AI and Skills Readiness: Workforce Transformation in India's Manufacturing Regions, National Education Services.

Pasupuleti, MK 2025v, Healthcare AI Readiness Index: Benchmarking Indian States Against G20 and BRICS Pathways, National Education Services.

Pasupuleti, MK 2025w, Topology for Trustworthy AI: Graphs, Tensors, and Robust Learning Systems, National Education Services.

Pasupuleti, MK 2025x, 'Swarm Autonomy at Scale: Graph Neural Control for Industry 5.0 and Disaster Response', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 05, no. 12, pp. 121–148.

Pasupuleti, MK 2025y, Federated Biointelligence: Privacy-Preserving Multi-Omics and Medical Imaging AI, National Education Services.

Pasupuleti, MK 2025z, Earth TwinOps: Geospatial AI and Quantum Optimization for Resilient Cities, National Education Services.

Pasupuleti, MK 2025aa, Net-Zero AI Infrastructure: Sustainable Computing from Cloud to Edge, National Education Services.

Pasupuleti, MK 2025ab, Quantum–AI Advantage: Hybrid Intelligence for National-Scale Systems, National Education Services.

eSign
Signed by: MURALIKRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication| Copyright© 2026
Location: Delhi, India
Date: Jan 30 2026 (05:00 PM).

Pasupuleti, MK 2025af, 'Multi-Omics for Transforming Healthcare: Integrating Genomics, Proteomics, and AI for Tailored Treatments', in Multi-Omics for Transforming Healthcare, National Education Services, pp. 71–82.

Pasupuleti, MK 2025ag, 'Sustainable Materials for Clean Energy: Nanomaterials, Solar Cells, and Hydrogen Innovation', in Sustainable Materials for Clean Energy, National Education Services, pp. 96–110.

Krishna Pasupuleti, M 2025ah, 'AI-Enabled Multimodal Biometrics: Advancing Security with Facial, Voice, and Behavioral Integration', in AI-Enabled Multimodal Biometrics, National Education Services, pp. 1–9.