WEEK 3

ARGO MODULE 1 | DATA SECURITY, PRIVACY, AND ETHICS IN LMIC HEALTH SYSTEMS

Ethics, Data Sovereignty, and Governance

The questions regulation alone cannot answer: power, benefit, and who decides.

Topic 1
Algorithmic bias and AI risk

Topic 2
Data sovereignty and donor dynamics

Topic 3
Community data ownership

Lecture notes | Live session: The Ethics Case

FRAMING

Why this week matters

- Weeks 1 and 2 built the threat model and the legal map. Compliance with both can still leave a data practice unjust.
- Who owns the health data generated by donor-funded programs in sovereign nations?
- When a model trained on African clinical data is commercialized in the Global North, who benefits and who bears the risk?
- When a state builds biometric health identity, what happens to people with rational reasons to avoid identification?

The point: The week does not pretend these questions have clean answers. It supplies frameworks for reasoning through them rigorously, and defensibly.

TOPIC 1

How bias enters health AI

- **Training data skew:** Models built predominantly on high-income-country populations meet different disease prevalence, demographics, and clinical presentation on deployment.
- **Label and measurement bias:** When cost or access history stands in for need, the model learns the inequity and reproduces it.
- **Deployment mismatch:** Lighting, connectivity, language, and workflow differ from the lab; performance quietly degrades where oversight is thinnest.
- **Feedback loops:** Resource-allocation models shape the data they are retrained on, hardening the pattern.

The point: Most LMIC jurisdictions have no regulatory framework for health AI. The governance burden falls on procurement and deployment decisions.

TOPIC 1

Documented cases

- **Dermatology AI:** Performance disparities across skin tones are documented in the literature; training-set composition is the mechanism.
- **Retinal screening in Thailand:** A high-accuracy model met clinic reality: image quality thresholds, connectivity, and workflow produced rejected scans and frustrated nurses.
- **Maternal risk scoring in India:** Predictive analytics for maternal mortality risk raises the question of what a risk score encodes, and what acting on it redistributes.

The evaluation question

Not “is the model accurate?” but “accurate for whom, measured where, and validated against which population?” Procurement that cannot ask this question cannot protect patients.

TOPIC 2

Data sovereignty operates on three levels

- **National:** Localization requirements, jurisdiction over data, and the state's claim to govern information about its population.
- **Community and Indigenous:** The authority of the people the data describes, which national control does not automatically deliver.
- **Operational:** Who holds the keys, the analytic capacity, and the contract terms. Server location is not the same thing as control.

The point: A dataset hosted in-country under a license the ministry cannot renegotiate is localized, not sovereign.

TOPIC 2

Donor dynamics: reading the fine print

- In donor-funded systems, technical infrastructure, analysis capacity, and sometimes the data itself sit outside the country.
- Licensing agreements embed control: access terms, export rights, termination clauses, and who owns derivatives.
- Open-data mandates from global health funders now collide with national data sovereignty positions; both sides have real arguments.
- Reporting burdens shape what gets measured, which quietly shapes what the health system attends to.

Applied: Exercise: dissect a real data-sharing agreement between a ministry of health and an international implementing partner, marking each provision as protecting or undermining sovereignty.

TOPIC 2

The CARE Principles for Indigenous Data Governance

Principle	What it asks	Applied to a national health dataset
Collective benefit	Data ecosystems enable the people the data describes to benefit	Benefit-sharing terms in any secondary-use or AI-training agreement
Authority to control	The described community holds decision rights	Community representation with a vote, not a seat, in data governance
Responsibility	Those working with the data are accountable to the community	Named accountability for misuse; audit rights that run to communities
Ethics	Rights and wellbeing govern across the data lifecycle	Harms assessment before secondary use, not after publication

Note: CARE was designed to complement FAIR, shifting the question from whether data can flow to who decides and who benefits.

TOPIC 3

OCAP and community data ownership

- **OCAP:** Ownership, Control, Access, Possession, developed by the First Nations Information Governance Centre.
- The transferable insight: possession is a governance instrument. Holding the data is a form of authority that contracts alone do not replace.
- Transfer the lesson with care for context; OCAP belongs to a specific history and is not a template to copy-paste.
- Related models: community data governance piloted by Last Mile Health in Liberia, and the emerging discourse on data trusts for health data in the Global South.

TOPIC 3

From principles to practice

- Community representation in data governance committees, with decision rights over secondary use, including AI training.
- Benefit-sharing written into data-sharing agreements, not promised alongside them.
- Local hosting and capacity clauses that build the ability to say no, not just the right to.
- Accountability mechanisms for misuse that a community can actually invoke.

Honest tension: The policy-implementation gap runs through this week too: laws without enforcement budgets, DPO roles assigned but not resourced, donor timelines that outpace governance timelines. Naming the gap is a design input for recommendations, not cynicism.

LIVE SESSION (90 MINUTES)

The Ethics Case

- **The setup:** A large global health organization funded a digital platform in three East African countries, collecting granular individual-level maternal and child health data.
- **The proposal:** Make a de-identified version of the dataset available to a machine learning consortium, including a private health-technology company, to build a maternal mortality risk model.
- **The friction:** Three countries, three legal frameworks, three positions on sovereignty, three levels of capacity to evaluate the proposal.

The point: De-identified is doing heavy lifting in that sentence. Interrogate it.

LIVE SESSION

The stakeholder map

Stakeholder	Interests	Power and vulnerability
Patients who generated the data	Care quality; privacy; not being experimented on invisibly	Least power; highest exposure; no seat at the table by default
Community health workers	Tools that work; not becoming the extraction layer	Frontline knowledge; little contractual leverage
National governments	Sovereignty; capacity; health outcomes; leverage with funders	Regulatory authority; dependent on external financing
The funding organization	Impact; reputation; the value of the asset it financed	Controls infrastructure and terms
The ML consortium	Data access; publishable and commercial results	Technical capacity; needs legitimacy

FINAL ASSESSMENT

Ethics Case Analysis (60% of module grade)

- **Scope:** 3,000 to 3,500 words. Use the live session case or one of the provided alternatives. Due end of Week 4.
- **Required elements:** Stakeholder mapping (interests, power, vulnerability); at least two ethical frameworks, drawn from principlism, Ubuntu ethics, feminist ethics of care, the CARE Principles, or other established frameworks; regulatory analysis and its limits; a clear recommendation with justification; an implementation plan honest about constraints.
- **Grading:** Quality of ethical reasoning, sophistication of stakeholder analysis, and grounding in the realities of the setting rather than abstract principle.

Use the cohort: Post your stakeholder map to forum thread 3.3 for peer review before drafting. The best comment names the stakeholder you missed.

SYNTHESIS

What the module adds up to

- Week 1 built a threat model that respects real constraints.
- Week 2 mapped two legal layers and the clocks they run on.
- Week 3 asks the questions law cannot settle: who decides, and who benefits.
- A governance recommendation that ignores any of the three is incomplete, and the final assessment is built to detect the gap.

Closing forum work

Thread 3.1: apply one CARE principle to a dataset you manage. Thread 3.2: draft the five-sentence reply to a funder demanding record-level export. Thread 3.3: post your capstone stakeholder map for peer comment.